



CARDIFF UNIVERSITY
SCHOOL OF EARTH AND OCEAN SCIENCES

PRIFYSGOL CAERDYDD
YSGOL GWYDDORAU'R DDAEAR A'R MÔR

**APPLYING MULTI-RESOLUTION NUMERICAL METHODS TO
GEODYNAMICS**

David RHODRI Davies

SUPERVISORS: J. H. Davies, O. Hassan, K. Morgan and P. Nithiarasu

*Thesis submitted to Cardiff University, UK
in partial fulfilment for the degree of Doctor of Philosophy*

May 2008

UMI Number: U585288

All rights reserved

INFORMATION TO ALL USERS

The quality of this reproduction is dependent upon the quality of the copy submitted.

In the unlikely event that the author did not send a complete manuscript and there are missing pages, these will be noted. Also, if material had to be removed, a note will indicate the deletion.



UMI U585288

Published by ProQuest LLC 2013. Copyright in the Dissertation held by the Author.
Microform Edition © ProQuest LLC.

All rights reserved. This work is protected against
unauthorized copying under Title 17, United States Code.



ProQuest LLC
789 East Eisenhower Parkway
P.O. Box 1346
Ann Arbor, MI 48106-1346

Abstract

Computational models yield inaccurate results if the underlying numerical grid fails to provide the necessary resolution to capture a simulation's important features. For the large-scale problems regularly encountered in geodynamics, inadequate grid resolution is a major concern. The majority of models involve multi-scale dynamics, being characterized by fine-scale upwelling and downwelling activity in a more passive, large-scale background flow. Such configurations, when coupled to the complex geometries involved, present a serious challenge for computational methods. Current techniques are unable to resolve localized features and, hence, such models cannot be solved efficiently. This thesis demonstrates, through a series of papers and closely-coupled appendices, how multi-resolution finite-element methods from the forefront of computational engineering can provide a means to address these issues.

The problems examined achieve multi-resolution through one of two methods. In two-dimensions (2-D), automatic, unstructured mesh refinement procedures are utilized. Such methods improve the solution quality of convection dominated problems by adapting the grid automatically around regions of high solution gradient, yielding enhanced resolution of the associated flow features. Thermal and thermo-chemical validation tests illustrate that the technique is robust and highly successful, improving solution accuracy whilst increasing computational efficiency. These points are reinforced when the technique is applied to geophysical simulations of mid-ocean ridge and subduction zone magmatism. To date, successful goal-orientated/error-guided grid adaptation techniques have not been utilized within the field of geodynamics. The work included herein is therefore the first geodynamical application of such methods.

In view of the existing three-dimensional (3-D) spherical mantle dynamics codes, which are built upon a quasi-uniform discretization of the sphere and closely coupled structured grid solution strategies, the unstructured techniques utilized in 2-D would throw away the regular grid and, with it, the major benefits of the current solution algorithms. Alternative avenues towards multi-resolution must therefore be sought. A non-uniform structured method that produces similar advantages to unstructured grids is introduced here, in the context of the pre-existing 3-D spherical mantle dynamics code, TERRA. The method, based upon the multigrid refinement techniques employed in the field of computational engineering, is used to refine and solve on a radially non-uniform grid. It maintains the key benefits of TERRA's current configuration, whilst also overcoming many of its limitations. Highly efficient solutions to non-uniform problems are obtained. The scheme is highly resourceful in terms RAM, meaning that one can attempt calculations that would otherwise be impractical. In addition, the solution algorithm reduces the CPU-time needed to solve a given problem. Validation tests illustrate that the approach is accurate and robust. Furthermore, by being conceptually simple and straightforward to implement, the method negates the need to reformulate large sections of code.

The technique is applied to highly advanced 3-D spherical mantle convection models. Due to its resourcefulness in terms of RAM, the modified code allows one to efficiently resolve thermal boundary layers at the dynamical regime of Earth's mantle. The simulations presented

are therefore at superior vigor to the highest attained, to date, in 3-D spherical geometry, achieving Rayleigh numbers of order 10^9 . Upwelling structures are examined, focussing upon the nature of deep mantle plumes. Previous studies have shown long-lived, anchored, coherent upwelling plumes to be a feature of low to moderate vigor convection. Since more vigorous convection traditionally shows greater time-dependence, the fixity of upwellings would not logically be expected for non-layered convection at higher vigors. However, such configurations have recently been observed. With hot-spots widely-regarded as the surface expression of deep mantle plumes, it is of great importance to ascertain whether or not these conclusions are valid at the dynamical regime of Earth's mantle. Results demonstrate that at these high vigors, steady plumes do arise. However, they do not dominate the planform as in lower vigor cases; they coexist with mobile and ephemeral plumes and display ranging characteristics, which are consistent with hot-spot observations on Earth. Those plumes that do remain steady alter in intensity throughout the simulation, strengthening and weakening over time. Such behavior is caused by an irregular supply of cold material to the core-mantle boundary region, suggesting that subducting slabs are partially responsible for episodic plume magmatism on Earth.

With this in mind, the influence of the upper boundary condition upon the planform of mantle convection is further examined. With the modified code, the CPU-time needed to solve a given problem is reduced and, hence, several simulations can be run efficiently, allowing a relatively rapid parameter space mapping of various upper boundary conditions. Results, in accordance with the investigations on upwelling structures, demonstrate that the surface exerts a profound control upon internal dynamics, manifesting itself not only in convective structures, but also in thermal profiles, Nusselt numbers and velocity patterns. Since the majority of geodynamical simulations incorporate a surface condition that is not at all representative of Earth, this is a worrying, yet important conclusion. By failing to address the surface appropriately, geodynamical models, regardless of their sophistication, cannot be truly applicable to Earth.

In summary, the techniques developed herein, in both 2- and 3-D, are extremely practical and highly efficient, yielding significant advantages for geodynamical simulations. Indeed, they allow one to solve problems that would otherwise be unfeasible.

Contents

Title Page	i
Abstract	iii
Table of Contents	v
Acknowledgments	ix
Dedication	x
1 Introduction	1
1.1 Motivation & Background	1
1.2 A Geophysical Primer	3
1.2.1 Earth Was Born	3
1.2.2 Earth's Structure	6
1.2.3 Plate Tectonics	12
1.2.4 Summary	15
1.3 Computational Background	15
1.3.1 Geodynamical Modeling	17
1.3.2 Summary	19
1.4 Thesis Outline	20
2 Adaptive Finite Element Methods for Two-Dimensional Infinite Prandtl Number Thermal & Thermo-Chemical Convection	22
2.1 Abstract	22
2.2 Introduction	23
2.3 Problem Description	27
2.3.1 Case 1: Thermal Convection in a Square Cavity	29
2.3.2 Case 2: Entrainment of a Thin Dense Layer through Thermo-chemical Convection	30
2.4 Adaptive Methodology	32
2.4.1 Remeshing Procedure for Case 1 - Steady-State Adaptivity	35
2.4.2 Remeshing Procedure for Case 2 - Temporal Adaptivity	36
2.5 Results & Discussion	37
2.5.1 Thermal Convection in a Square Cavity	37
2.5.2 Thermo-chemical Convection	47
2.6 Conclusions	55

2.6.1	Applicability to Thermal Convection	56
2.6.2	Applicability to Thermo-chemical Convection	56
3	Adaptive Finite Element Methods in Geodynamics; Convection Dominated Mid-Ocean Ridge and Subduction Zone Simulations	58
3.1	Abstract	58
3.2	Introduction	59
3.3	Methodology	62
3.3.1	Governing Equations and Solution Procedure	62
3.3.2	Adaptive Strategies	63
3.4	Geodynamical Applications	65
3.4.1	Mid-Ocean Ridge Magmatism	65
3.4.2	Subduction Zone Magmatism	73
3.5	Conclusions	82
4	Towards Automatic Multi-resolution in 3-D Spherical Models of Mantle Convection	83
4.1	Abstract	83
4.2	Introduction	84
4.2.1	Motivation and Background	84
4.2.2	The Need for Multi-Resolution	86
4.3	Methodology	88
4.3.1	Multigrid	88
4.3.2	The Full Approximation Storage Scheme	91
4.3.3	The Adaptive Multigrid	93
4.3.4	Implementation and Solution Procedures	97
4.4	Implementation within TERRA	100
4.4.1	Model Discretization	103
4.4.2	Converting TERRA's Multigrid From CS To FAS	105
4.4.3	Adaptive Solution Strategy for Force-Balance Equation	108
4.4.4	The Multi-Level Energy Equation	111
4.5	Validation of TERRA and Radial Refinement Scheme	114
4.5.1	Thermal Amplitude Tests	115
4.5.2	Stokes Flow	116
4.6	Computational Efficiency	118
4.7	Conclusions	124
5	Mobile and Steady Plumes Coexist at Earth-like Vigor	126
5.1	Abstract	126
5.2	Introduction	127
5.3	Methodology	129
5.3.1	Numerical Methodology	129
5.3.2	Model Configuration	131
5.4	Results & Discussion	134
5.4.1	Isoviscous Convection	134

5.4.2	Layered Viscosity Convection	137
5.5	Discussion	145
5.6	Conclusions	150
6	The Importance of Earth's Surface Boundary Condition to Mantle Convection Simulations	152
6.1	Abstract	152
6.2	Introduction	153
6.3	Methodology	156
6.3.1	Numerical Methodology	156
6.3.2	Model Configuration	157
6.4	Results & Discussion	162
6.5	Conclusions	168
7	Conclusions	170
A	The Finite Element Method	193
A.1	Introduction	193
A.2	Overview	194
A.3	Equation Discretization	196
A.3.1	Mathematical Preliminaries	196
A.3.2	Trial Solutions and Weighting Functions	199
A.3.3	Strong and Weak Forms of a Boundary Value Problem	200
A.3.4	Galerkin's Approximation Method	203
A.4	Elements and Shape/Interpolation Functions	207
A.4.1	One-dimensional Linear Elements	208
A.4.2	One-dimensional Quadratic Element	211
A.4.3	Two-dimensional Linear Triangular Elements	215
A.4.4	Requirements for Interpolation Functions	218
A.5	Spatial Discretization	219
B	CONMAN	223
B.1	Introduction	223
B.2	Equations and implementation	223
B.3	The Momentum Equation	225
B.3.1	Preliminaries	225
B.3.2	Prescribed Data and Formal Statement of the Problem	226
B.3.3	Penalty Function Formulation	227
B.3.4	Weak Formulation	228
B.3.5	Galerkin Formulation	230
B.3.6	Matrix Problem	230
B.4	The Energy Equation	233
B.4.1	Predictor-Multicorrector Algorithm	234

C	Two Dimensional Mesh Generation and Adaptive Remeshing Procedure	237
C.1	Mesh Generation	237
C.1.1	Generation of the Initial Mesh	238
C.1.2	Mesh Control: The Background Mesh	239
C.1.3	Boundary Representation	240
C.1.4	Triangle Element Generation	241
C.1.5	Quadrilateral Element Generation using an Existing Triangular Mesh	242
C.1.6	Mesh Quality Enhancement	243
C.2	The Error Indicator and Adaptive Strategy	244
C.2.1	The Error Indicator	246
C.2.2	Adaptive Remeshing	247
D	Calculation of Nodal Gradients and Curvatures	250
E	TERRA	252
E.1	Overview & History	252
E.2	The Governing Equations	253
E.3	The Computational Grid	257
E.3.1	Indexing Conventions Employed by Terra	259
E.4	The Galerkin Finite Element Formulation	261
E.4.1	Matrix Equation Derivation	263
E.5	Computational Operators	266
E.5.1	The Finite Element Shape Functions	266
E.6	Solution Strategy for Force-Balance Equation	270
E.6.1	A Pressure Correction Approach	270
E.6.2	Matrix Dependent Transfer	274
E.6.3	Coarse Grid Operators	274
E.6.4	Relaxation Method	275
E.7	The Energy Equation and Time Integration	276
E.7.1	Cell-Wall Advection	276
E.7.2	Thermal Conduction	278
E.7.3	Time Integration	279
E.8	Parallelization	280
E.8.1	Overview	280
E.8.2	Technicalities	280
E.9	Visualization	282
F	Multigrid	283
F.1	Introduction	283
F.2	The Residual Equation	283
F.3	Relaxation Methods	284
F.4	Coarse Grid Correction	286
F.4.1	V-Cycle Scheme	289
F.5	The Full Approximation Storage Scheme	289
F.6	Conclusions	291

Acknowledgments

Completing this doctoral work has been a wonderful and often overwhelming experience. It is hard to tell whether getting to grips with the numerical techniques has been the real learning experience, or tackling problems such as how to write a paper, give a coherent talk, work in a group, code intelligently, recover a crashed hard drive, stay up until the birds start singing etc...

Either way, it has been a great experience. I have been privileged to have undoubtedly the most intuitive, smart and supportive advisors anyone could ask for, namely Huw Davies at Cardiff and Oubay Hassan, Ken Morgan and Perumal Nithiarasu at Swansea. Over the past three years I have been stimulated and excited by their constant flow of good ideas. They have truly challenged me to produce a well thought out piece of research that I am proud to present. I wish to sincerely thank them all for their support. They are individuals and friends that I will always admire.

I recognize the contributions of both Scott King and John Baumgardner to this thesis. Scott supplied the two-dimensional geodynamical code CONMAN, while John provided the three-dimensional code TERRA. Without these codes, the majority of results presented herein would have been unattainable. The simulations presented were executed on several large computing clusters, namely the IBM Deep Computing facility at Swansea University, the Helix (SRIF funding, HEFCW) and Cubric facilities at Cardiff University and, HECTOR, the current British supercomputer at EPCC, through the Mineral Physics Consortium at the Natural Environmental Research Council (NERC). Notes of thanks are given to Ian Merrick, Diane Cook, Jason Jones, Andrew Austin, Alun Rogers and Spiro Stathakis for computational support.

The staff and students at Cardiff and Swansea Universities are acknowledged. Both institutions have provided excellent research environments. Ben Evans, David Oldham and Richard Davies are thanked for helpful and informative discussion. Martin Wolstencroft, who offered significant help with proof-reading, visualization etc... during the final weeks before submission, is also gratefully acknowledged. In addition, the Computational Infrastructure in Geodynamics (CIG) team at CALTECH, particularly Mike Gurnis, are thanked for their advice and support over the past three years. Peter Bunge, Tim Phillips and David Rickard are acknowledged for a detailed and thorough examination of this thesis.

I gratefully acknowledge the financial support of NERC (NER/S/E/2004/12725) as part of the Environmental Mathematics and Statistics Scheme.

Finally, I wish to take this opportunity to thank my incredible family, my great friends and my wonderful girlfriend, for their constant, unwavering support and for continually encouraging me to do my best. I reserve special thanks to my parents, who have been with me through every peak and trough *en route* to submission. I will be eternally grateful to you both for everything.

I dedicate this thesis to my family and friends, particularly my parents, Neville and Jennifer. Without you, this work wouldn't have been possible.

Chapter 1

Introduction

1.1 Motivation & Background

Earth's present-day mantle is obscured from direct view. However, thanks to disciplines such as seismology, geodesy, geochemistry and mineral physics, a wealth of information has accumulated regarding its structure, composition and internal dynamics (e.g. Richards & Hager 1984, Allegre et al. 1996, van der Hilst et al. 1997, Oganov & Ono 2004). This information has been supplemented by numerical simulations of the mantle's convective flow field, yielding valuable insights into the temporal and spatial evolution of mantle flow under a range of rheologies and boundary conditions (e.g. Bunge et al. 1996, Zhong et al. 2000, Lowman et al. 2001). Such simulations have become an essential component in our attempts to understand the dynamics of Earth's mantle, however, for various reasons, which are discussed later, mantle convection codes are difficult to formulate and validate. Furthermore, the scale and geometry of the problems under investigation mean that, regardless of their numerical efficiency, even the simplest simulations are computationally highly demanding. As a consequence, despite the massively parallel computational architectures available to us today, it is often impossible to resolve thermal and chemical boundary layers

at the dynamical regime of Earth's mantle. This task becomes even more difficult for 'early Earth' simulations, likely to be characterized by more vigorous convection, which is a major concern.

It is broadly accepted that mantle convection involves fine scale upwelling and downwelling activity in a more passive, large-scale background flow. Such narrow, high gradient regions present a serious challenge for computational methods; their location and extent is often undeterminable and, even if they are identified, current methods are unable to resolve localized features. The primary motivation behind this thesis therefore is:

To develop computationally efficient, multi-resolution numerical techniques, which accurately solve such 'multi-scale' flows, thereby yielding a new class of geodynamical model.

Such techniques provide an avenue towards the secondary objective of this investigation:

To utilize these advanced models, in both two- and three- dimensional mantle convection and circulation simulations, to advance our understanding of a variety of geophysical processes.

The thesis incorporates both sophisticated numerical techniques and cutting-edge geophysical application. Such a multi-disciplinary project necessitates that a broad background to both geophysical and computational aspects be provided, allowing the reader to comprehend unfamiliar subject matter. Accordingly, section 1.2 provides a general synopsis of our current understanding of Earth's history, along with an overview of the structure and composition of Earth's interior, focussing upon the dynamics of mantle convection. This is intended to provide the more computationally inclined reader with a background to the geophysical problems under investigation. Likewise, an introduction to computational

fluid dynamics, numerical modeling and the finite element method is provided in section 1.3, giving those with more of a geophysical background a handle on the methodological sections included herein. Section 1.4 provides an outline of the thesis' content. Excluding this chapter and the general conclusions of Chapter 7, each chapter is a stand-alone paper, with successive papers representing a progression towards the objectives outlined above. Supplementary information is included in the form of various appendices.

1.2 A Geophysical Primer

Geophysics, the physics of the Earth, is a huge subject encompassing studies of space and the atmosphere, of the oceans and of the interior of the planet. The aim of this section is to convey, in a straightforward manner, what is known of the structure and dynamics of the solid Earth, focussing upon the mantle. It is designed to serve as an introduction to the areas of geophysics, relevant to this thesis, for engineers and mathematicians. Small sections are also included summarizing recent geophysical particulars that will be of interest to a more general reader.

1.2.1 Earth Was Born

Around 5 billion years ago, somewhere in our galaxy, a supernova exploded, pushing large amounts of heavy-element wreckage into a nearby cloud of hydrogen gas and interstellar dust. The mixture grew hot, compressing under its own gravity and, at its center a new star began to form. The stars' gravity caused the remaining gas cloud to be compressed into a white-hot rotating disk. That new star became our Sun, and the glowing disk gave rise to Earth and its sister planets.

While the Sun grew in size and energy, beginning to ignite its nuclear fires, the hot disk




Figure 1.1: The theory of planetary formation (not drawn to scale, Pluto not shown). (a)/(b) The solar nebula contracts and flattens into a spinning disk. The large blob in the center will later become the Sun. Smaller blobs in the outer regions may become planets. (c) Dust grains act as condensation nuclei, forming clumps of matter that collide, stick together, and grow into planetesimals. (d) Strong winds from the still-forming Sun expel the nebular gas. (e) Planetesimals continue to collide and grow. (f) Over the course of a hundred million years or so, planetesimals form a few large planets that travel in roughly circular orbits. Reproduced from <http://physics.fortlewis.edu>.

slowly cooled. This took millions of years. During that time, the components of the disk began to freeze out into small dust-size grains. Heavier atoms, such as iron, magnesium and silicon, were able to condense closer to the sun. Lighter elements, like hydrogen, could only condense in more distal parts of the disk, giving rise to variations that would later lead to the formation of inner ‘rocky’ planets and outer ‘gas giants’. Slowly these grains settled together and collected into clumps, then chunks, then boulders and finally bodies large enough to exert significant gravity - *planetesimals*.

As time went by, planetesimals grew by collision with other bodies and, as their mass grew larger, so did the energies involved. By the time they had reached 100 km or so in size, planetesimal collisions had the potential to cause outright melting and vaporization. Consequently, the materials, which we can confidently call rocks and iron metal, began to undergo segregation. The dense iron settled in the center and the lighter rock separated into a mantle around the iron. Planetologists call this *differentiation*. At around this time, the Sun ignited. The process of ignition was energetic enough to blow away most of the gaseous part of the protoplanetary disk. The chunks, boulders, and planetesimals left behind continued to collect into a handful of large, stable bodies in well-spaced orbits, surrounded by significant quantities of smaller debris. Earth was the third one of these, counting outward from the Sun (e.g. Safronov 1969, Wetherill & Stewart 1989, Weidenschilling et al. 1997, Chambers 2001). See Lissauer (1993) for a detailed review.

Since then, our planet has evolved from a gravitationally accumulated ball of debris into the world as we know it. The geological record illustrates that the oldest surviving rocks on Earth were formed ≈ 600 million years after planetary formation (≈ 3.9 billion years ago), although evidence from Zircon crystals point towards rock-forming activity as far back as ≈ 4.4 billion years ago (Wilde et al. 2001). These lines of evidence suggest that there were




Figure 1.2: A simplified representation of Earth's interior, believed to consist of several concentric spheres: an outer solid crust, averaging 7-8 km thickness beneath oceans and ≈ 35 km beneath continents; a mainly solid mantle, extending to a depth of 2900 km; an outer core of liquid iron, extending to a depth of 5150 km; and an inner core of solid iron, with a radius of about 1220 km. Modified from <http://observatory.ou.edu/EarthInterior.jpg>.

volcanoes, continents, oceans and life on Earth in those days. While the eons that followed manufactured significant developments and considerable changes, Earth had taken on its basic structure long before.

1.2.2 Earth's Structure

The major divisions of Earth's interior are shown in Figure 1.2. The rocks exposed at Earth's surface form part of the *crust*. This crustal layer, which is rich in silica, was identified by Milne (1906) and Rutherford (1907). On average, it has a thickness of 35 km beneath the continents and 7-8 km beneath the oceans. Below this thin crust lies the *mantle*, which extends down some 2900 km to Earth's central *core*. The mantle is both physically and chemically distinct from the crust, being rich in magnesium silicates. Indeed, the crust is derived from the mantle over time through a series of melting and

reworking processes. Earth's core was discovered by Oldham (1906) and correctly delineated by Gutenberg (1913) from studies of Earthquake data. The core is totally different, both physically and chemically, from the crust and mantle. It is predominantly composed of iron with lesser amounts of lighter elements. Jeffreys (1926) established that the core was fluid and, subsequently, using seismic data, Lehmann (1936) demonstrated that Earth has an *inner core* within the liquid outer core. The inner core is solid and grows as Earth cools.

The Crust

In terms of its mass, the crust is a negligible part of Earth, comprising less than 0.5% of the total; however, its significance is far greater than is suggested by this figure. It is the end product of mantle differentiation, yielding valuable information on the history of mantle convection. The crust also contains the most extreme lateral heterogeneities on Earth and a substantial fraction of Earth's incompatible elements (those that prefer to move into the liquid phase during melting), including the heat-producing elements uranium, thorium and potassium, directly influencing the style of mantle convection and Earth's thermal evolution.

Continental and oceanic crust is bimodal in nearly every respect: thickness, age, density and composition (e.g. Schubert et al. 2001). On average, the continental crust is old (mean age $\approx 2,000$ Myr, where Myr represents one million years) and has resided on Earth's surface long enough to have experienced a sequence of collision events, sedimentation, episodes of loading and unloading, intrusions of melts from below, infiltration of metamorphic fluids, erosion and alterations in its thermal regime. As a result, the continental crust varies greatly in the variety of its igneous, sedimentary and metamorphic rocks. On the other hand, the oceanic crust is relatively young (mean age ≈ 80 Myr), created continuously at spreading centers (the present rate of creation being $2.8 \text{ km}^2\text{yr}^{-1}$) and subducted back into the mantle at approximately the same rate (Schubert et al. 2001). It is formed almost entirely by partial

melting of the upper mantle, resulting in a remarkably homogeneous product, mid-ocean ridge basalt (MORB).

The Mantle

The mantle contains 84% of Earth's volume and 68% of its mass, but because it is separated from direct observation by the thin crust, there are many unsolved problems. Nevertheless, we know enough to give a broad picture of its structure and composition, which is the principal aim of this section. Mantle convection establishes one of the longest time scales of our planet. Though solid, the mantle is deforming slowly by a process of viscous creep and, while sluggish in human terms, the rate of this subsolidus convection is remarkable by any standard. Indeed, it is estimated that the mantle's Rayleigh number is of order 10^9 (Davies 1999, Schubert et al. 2001), generating flow velocities of 2–10 cm yr⁻¹. Plate tectonics is the prime surface expression of this convection (e.g. Davies & Richards 1992), although, ultimately, all large-scale endogenic geological activity and dynamics of the planet, such as mountain building and continental drift, involve the release of potential energy in the mantle.

Convection within Earth's present-day mantle is becoming increasingly well understood as the superposition of two main modes (Davies 1999). The dominant plate-scale flow is driven mainly by the gravitational instability of the cold, stiff upper thermal boundary layer or lithosphere. Plumes represent the second, subsidiary mode of mantle convection, arising from instabilities in a bottom-heated lower thermal boundary layer. The overall dynamics of mantle convection, however, remains one of the most controversial issues in Earth sciences.

Seismological images reveal three distinct changes in velocity structure, at depths of 410-, 660- and 2700-km. These are best explained by mineral phase transformations, with the

third possibly incorporating a change in chemical composition and thermal structure (Murakami et al. 2004, Oganov & Ono 2004, Tsuchiya et al. 2004). Recent debate has focused on whether Earth's upper and lower mantles are mixing. Several compelling geochemical arguments, described further below, require at least two distinct mantle reservoirs that have preserved chemical heterogeneity for billions of years. The boundary between the two reservoirs has usually been placed at the 660-km seismic discontinuity, an endothermic phase change which may act to impede flow.

On the basis of strong mass-balance arguments, geochemists have, in general, been strongly supportive of separate convective regimes with little or no mass flux between them. Together with the less radiogenic character of helium in most ocean-island basalts (OIBs) with respect to MORBs (O'Nions & Oxburgh 1983), the apparent imbalance between the terrestrial budgets of ^{40}K and ^{40}Ar was long considered the strongest evidence that the lower mantle is not degassed (Allegre et al. 1996). The terrestrial imbalance between heat production and surface heat flow on the one hand and the isotopic budget of radiogenic elements in chondritic Earth on the other, also seems to require the presence of substantial mantle volumes that are not sampled by MORB (e.g. Boyet & Carlson 2006). As noted previously, the sharp, global seismic discontinuity at a depth of 660-km was initially chosen as a major hindrance to whole-mantle convection on the basis that very few earthquakes seemed to originate from below this depth and, therefore, that lithospheric plates do not penetrate it. An additional justification of layering at this depth was that subtracting the inventory of incompatible elements of the crust and the upper mantle from the inventory of the bulk silicate Earth (BSE) leaves a nearly chondritic distribution of the lithophilic refractory elements in the lower mantle (Albarede & van der Hilst 2002). Even the terrestrial Nd and Sr isotopic budgets could be interpreted as being consistent with a lower mantle consisting of primitive Sm/Nd and Rb/Sr ratios (Depaolo & Wasserburg 1976, Allegre et al. 1980).




Figure 1.3: Seismic tomographers feel that 660 km discontinuity can no longer be defended as a barrier to flow. This is illustrated here, in a vertical mantle section from Earth's surface to the core-mantle boundary, across Indonesia. Dashed lines are drawn at depths of 410-, 660- and 1700-km respectively, and subducted slabs can clearly be seen passing through 660 km. Blue and red indicate fast and slow anomalies respectively. Adapted from Albarede & van der Hilst (2002).

In contrast, physical arguments raised by dynamicists (e.g. Hager 1984, Davies 1988, Tackley et al. 1994) and seismic imaging indicate that this discontinuity is by no means an absolute barrier to mantle flow. Two high-resolution seismic tomography models (Grand et al. 1997, van der Hilst et al. 1997) show fast slab-like anomalies extending down into the lower mantle, through the 660-km discontinuity (see Figure 1.3). These models display analogous features in well-resolved regions, with one study imaging the compressional (P-wave) and the other the shear wave (S-wave) structure, from different data sets using different methods. Results have been confirmed by more recent investigations (e.g. Bijwaard & Spakman 1999, Grand 2002, Becker & Boschi 2002).

It is clearly beyond the scope of this introduction to review the evidence supporting each point of view, however, tomographic images of cold slabs in the lower mantle, the displacements of the 410-km and 660-km discontinuities around subduction zones, and the occur-

rence of small-scale heterogeneities in the lower mantle all indicate that subducted material penetrates the deep mantle, implying whole-mantle convection. Although the simple model of two-layered convection appears, for the moment at least, to be out of favor, many of the arguments that support it are still difficult for a whole-mantle-convection model to explain. Several concepts have been proposed in recent years to resolve the apparent conflicts between the various observational constraints and theoretical interpretations. The suggestions include the presence of deeper layering (e.g. Kellogg et al. 1999), the preservation of highly viscous blobs and poor mixing (van Keken et al. 2002), a zoned, but not strictly layered mantle (Albarede & van der Hilst 2002), core mantle interaction (Brandon & Walker 2005), and strong temporal variations in mantle dynamics. Although these models generally appear to solve parts of the puzzle, as of yet, no single model appears to account for all major observations. This is an area of ongoing research (e.g. Parman 2007, Gonnermann & Mukhopadhyay 2007, Schmerr & Garnero 2007, Goes et al. 2008).

Earth's Core

Earth's core contains 16% of Earth's volume and 31% of its mass, with the inner core currently accounting for only $\approx 0.7\%$ and $\approx 1.7\%$ respectively. The core apparently lacks the complexity of the mantle and crust, though this apparent simplicity may be because of the difficulties involved in observing regions so deep inside Earth. However, it plays a key role in many geophysical processes - amongst other things, Earth's magnetic field is a result of motions in the fluid outer core.

Major advances towards a better understanding of the composition, structure and dynamics of Earth's core are to be gained only by a combination of experimental and theoretical techniques. There are many unanswered questions about the exact nature of the core and its interaction with the mantle. Nevertheless, we know enough to give a broad picture of

its structure and main compositional characteristics:

1. We understand from seismology that the core is divided into an outer liquid part and an inner solid part, with the outer core having a radius of 3485 km (Masters & Shearer 1990).
2. As Earth cools, the core solidifies, growing the solid inner core whose radius is currently about 1225 km.
3. Latent heat release and chemical differentiation at the inner core boundary drive convective motions in the liquid outer core, generating Earth's magnetic field.

Core composition however, remains poorly constrained and is a subject of lively, occasionally even acrimonious, debate. The major problem with theories of core composition is that they depend on theories of the origin of Earth and its chemical and thermal evolution, which are poorly understood. The work of Birch (1964) and others in the 1950s made it clear that the core consists mainly of iron with a minor fraction of nickel, but that unidentified light impurities must also be present, since the density of the core is significantly lower than that of iron: modern estimates (e.g. Stixrude et al. 1997) indicate that the outer core is lower in density by 5-10%, and the inner core by 2-3%. Cosmochemical abundances of the elements, combined with models of Earth's history, limit the possible impurities to a few candidates (Poirier 1994). The most often discussed are sulphur (Li & Agee 2001), oxygen (Alfe et al. 2000) and silicon (Gessmann et al. 2001).

1.2.3 Plate Tectonics

Earth has a cool and mechanically strong outermost shell, which is termed the *lithosphere*. It is of order 100 km thick and comprises the crust and uppermost mantle, being thinnest in oceanic regions and thicker in continental regions, where the depth of its base is poorly defined.

The *asthenosphere* is that part of the mantle immediately beneath the lithosphere. The high temperatures and pressures that exist at asthenospheric depth generate a region of low viscosity, where viscous flow occurs over geological timescales (millions of years, not seconds!). The basic concept of *plate tectonics* is that the lithosphere is divided into a small number of nearly rigid plates, which are moving over the asthenosphere. Most of the deformation which results from plate motion, such as stretching, folding or shearing, takes place along the edge, or boundary, of a plate. Deformation away from the boundary is not significant. Indeed, nearly all earthquakes, as well as most of Earth's volcanism, occur along plate boundaries. Plate boundaries are of three types:

1. *Divergent* - plates move away from each other. At such boundaries, new plate material, derived from the mantle, is added to the lithosphere. The divergent plate boundary is represented by the mid-ocean ridge system, along the axis of which new plate material is produced.
2. *Convergent* - plates move towards each other. Most such boundaries are represented by the oceanic-trench island-arc systems of subduction zones where one of the colliding plates descends into the mantle and is destroyed. Some convergent boundaries occur on land. Japan, the Aleutians and the Himalayas are the surface expression of convergent plate boundaries.
3. *Conservative* - lithosphere is neither created or destroyed. The plates move laterally relative to each other. These plate boundaries are represented by transform faults, of which the San Andreas Fault in California, USA, is a famous example.

There are seven main plates, of which the largest is the Pacific plate, and numerous smaller plates, such as Nazca, Cocos and Scotia (Figure 1.4). The theory of plate tectonics simply describes the interactions of the lithospheric plates and the consequences of these interactions.




Figure 1.4: The major tectonic plates, divergent, convergent and conservative boundaries. Note that this image has been reproduced from www.johomaps.com/world/worldtecton.html.

1.2.4 Summary

The synopsis presented above provides a solid basis to the areas of geophysics covered herein. As is clear from this short section, the study of Earth's interior is a wide-ranging, multidisciplinary area of research. In this thesis, the data and constraints provided by this multiplicity of sources is utilized to investigate, via numerical models, many of the fundamental processes that shape our planet. Attention is focussed upon the dynamics of whole-mantle convection. The complex relationship between the mantle and Earth's surface is also examined, particularly at divergent and convergent plate boundaries.

1.3 Computational Background

The goal of this section is to provide a background to the numerical tools utilized in this thesis. The advanced numerical techniques and computational algorithms developed herein are described in further detail where relevant. This section merely outlines the basic methodological concepts.

As stated in Section 1.2.2, Earth's mantle is solid. However, it deforms slowly over time and, consequently, motion within the mantle is described by the equations governing fluid mechanics. Broadly speaking, one has the following options when investigating fluid phenomena:

1. Taking measurements.
2. Use of analytical solutions.
3. Use of computer simulations.

The latter, commonly known as Computational Fluid Dynamics (CFD), generally involves finding the solution to the conservation equations of mass, momentum and energy in a fluid,

with very little simplification to these equations. The technique has rapidly established itself as the primary tool in fluid analysis, mainly because analytical solutions are often impossible to attain, whilst experimental methods are regularly limited for the following reasons:

1. It is often difficult to generate geometric similarity.
2. It can be impossible to obtain materials capable of replicating a fluids properties.
3. Running and manufacturing costs can be high. With these high costs, one might expect high accuracy. Often, this is not the case.

CFD, on the other hand, can provide:

1. A detailed understanding of the problem: CFD allows fluid properties and flow parameters, such as velocity, to be evaluated at any point in the flow. A more complete picture of flow phenomenon can therefore be derived, since data values exist everywhere and are not restricted to those areas instrumented during testing.
2. Rapid evaluation of design alternatives and optimization of performance: Once the CFD model has been built and validated, it is easy to change geometrical features and fluid parameters. In this way, a range of ‘what if’ scenarios can be quickly evaluated, either as a problem solving exercise or to optimize performance.
3. Reduction in testing costs: properly calibrated CFD models can reduce, or even eliminate, the need for expensive or large-scale test facilities. The investment in the relatively low cost, but powerful, computer hardware and CFD software can therefore provide significant savings in capital expenditure and development costs.
4. Reduction in failure risks: CFD can be used to investigate configurations that may be too large to test or that pose a significant safety risk.
5. Reduced time consumption: With the ability to re-use information generated at other stages of the process, in addition to rapid evaluation of design alternatives, CFD can be significantly less time consuming than experimental methods.

Due to the geometrical and rheological complexities inherent to convection within Earth’s mantle, CFD is one of the most suitable options available for examining these systems.

1.3.1 Geodynamical Modeling

Since the pioneering work of McKenzie et al. (1973, 1974), fluid dynamical simulations have stirred significant interest in the geodynamical community. As discussed in the previous section, such techniques provide a means to approximate, in an efficient manner, solutions to the governing equations of a specific problem. In geodynamics, the numerical methods often employed have been based upon the finite difference method, FDM (e.g. Bodri & Bodri 1978, Matyska & Yuen 2001) and the finite volume method, FVM (e.g. Tackley 1996, 1998, Ratcliff et al. 1998, Albers & Christensen 2001). The FDM is the oldest, centered upon the application of a local Taylor expansion to approximate the differential equations. The method uses a topologically square network of lines to construct the discretization of a differential equation. This, however, leads to a lack of flexibility and is the method's major bottleneck when handling complex geometries in multiple dimensions. This issue motivated the use of an integral form of the differential equations and subsequently the development of the FVM, which is similar to the FDM, in that values are calculated at discrete places on a meshed geometry. The FVM however is more flexible, since it can deal with non-uniform domain discretizations.

Over recent decades, another method, the Finite Element Method, FEM, which again is based upon an integral formulation, has become more established (e.g. Baumgardner 1985, Sidorin et al. 1999, Zhong et al. 2000). It is the FEM that is employed throughout this thesis. A brief introduction to the method is provided in the following section, however, since the approach is central to the investigations presented herein, a detailed account is provided in Appendix A. A reader unfamiliar with the technique should refer to this supplement before continuing.

The Finite Element Method (FEM)

The FEM considers that the solution region comprises many small, interconnected, sub-regions or elements and gives a piece-wise approximation to the governing equations. As a result, complex partial differential equations (PDEs) are reduced to either linear or non-linear simultaneous equations. Thus the finite element discretization procedure (i.e. dividing the domain into a number of smaller sub-domains) reduces the continuum problem, which has an infinite number of unknowns, to one with a finite number of unknowns at specified points, termed nodes. Since the FEM allows us to form the elements, or sub-regions, in an arbitrary sense, a close representation of the boundaries of complicated domains is possible.

The solution of a continuum problem by the FEM is approximated by the following step-by-step process (Lewis et al. 2004). This method is fairly general and is described further in Appendix A.

1. *Discretize the domain*: Divide the solution region into non-overlapping elements (see Figure 1.5). The finite element discretization allows a variety of element shapes, e.g. triangles and quadrilaterals in two-dimensions (2-D) and tetrahedra in three-dimensions (3-D). Each element is formed by the connection of a certain number of nodes.
2. *Select interpolation or shape functions*: Select the interpolation function to represent the variation of the field variable, e.g. temperature, over an element.
3. *Form element equations (formulation)*: Next, matrix equations are determined. These equations must express the properties of the individual elements.
4. *Assemble the element equations to obtain a system of simultaneous equations*: To find the properties of the overall system, one must assemble all individual element equations. The matrix equations of each element are combined in an appropriate way such that the resulting matrix represents the behavior of the entire computational




Figure 1.5: A typical quadrilateral finite element mesh. Elements, nodes and edges.

domain. Boundary conditions are incorporated after assemblage of the individual element contributions.

5. *Solve the system of equations:* The resulting set of algebraic equations is solved to obtain the nodal values of the field variable.
6. *Calculate the secondary quantities:* From the nodal values of the field variable, secondary quantities, e.g. the Nusselt number, are calculated.

More details on each of these points can be found in Appendix A.

1.3.2 Summary

Having introduced the aspects of geophysics and computational modeling relevant to the work included herein, the reader, regardless of intellectual expertise, should have ample background to engage in the remainder of this thesis. An outline of the overall structure of the thesis is next provided. This allows the reader to comprehend how computational

and geophysical aspects coalesce, or indeed, how the advanced computational techniques developed herein allow more challenging geophysical problems to be investigated.

1.4 Thesis Outline

The work presented falls naturally into two-sections. The first part covers 2-D models, while the second examines 3-D spherical simulations. Nonetheless, the second section provides a natural progression from the first.

The thesis begins with a paper that represents the first geodynamical exploitation of an advanced computational method, namely, adaptive mesh refinement. The technique provides a means to improve solution accuracy whilst increasing computational efficiency. The methodological concepts are introduced in detail and these are subsequently validated via thermal and thermo-chemical benchmark comparisons (Chapter 2). This methodology is then applied in 2-D mid-ocean ridge and subduction zone simulations (Chapter 3), illustrating the benefits to geophysical modeling.

The work is extended into 3-D spherical geometry in the second part of this thesis, albeit using different techniques. Based upon the adaptive multigrid methodologies employed in the field of computational engineering and the well-established mantle convection code TERRA (Baumgardner 1985), a novel procedure is introduced that allows one to refine and solve on a radially non-uniform grid. The major benefit to such a scheme is that complex problems can be solved efficiently, both in terms of RAM and CPU-time, while the methodology, being conceptually simple and straightforward to implement, negates the need to reformulate large sections of code.

A paper summarizing the technique, which we term ‘radial refinement’, and, subsequently,

its validation, is presented in Chapter 4. Two papers demonstrating the application of the method to highly advanced 3-D geophysical models are then presented. In Chapter 5, the nature of upwelling structures is investigated in simulations that replicate the dynamical regime of Earth's mantle, which, as noted previously, convects at a Rayleigh number of order 10^9 (Schubert et al. 2001). Were it not for the resourcefulness of the modified code such calculations would be implausible, as is demonstrated by the lack of current publications at these vigors. In Chapter 6, the influence of the upper boundary condition upon the planform of mantle convection is examined. The modified and highly efficient solution algorithm allows one to perform a rapid parameter space mapping, which is extremely beneficial. The thesis ends with a review of the principal results and a summary of the major conclusions (Chapter 7).

Chapter 2

Adaptive Finite Element Methods for Two-Dimensional Infinite Prandtl Number Thermal & Thermo-Chemical Convection

2.1 Abstract

An adaptive finite element procedure is presented for improving the quality of solutions to convection dominated problems in geodynamics. The method adapts the mesh automatically around regions of high solution gradient, yielding enhanced resolution of the associated flow features. The approach requires the coupling of an automatic mesh generator, a finite element flow solver and an error indicator. In this study, the procedure is implemented in conjunction with the well-known geodynamical finite element code, CONMAN. An unstructured quadrilateral mesh generator is utilized, with mesh adaptation accomplished through

regeneration. This regeneration employs information provided by an interpolation based local error indicator, obtained from the computed solution on an existing mesh. The technique is validated by solving thermal and thermo-chemical problems with well-established benchmark solutions. In a purely thermal context, results illustrate that the method is highly successful, improving solution accuracy whilst increasing computational efficiency. For thermo-chemical simulations the same conclusions can be drawn. However, results also demonstrate that the grid based methods employed for simulating the compositional field are not competitive with the other methods (tracer particle and marker chain) currently employed in this field, even at the higher spatial resolutions allowed by the adaptive grid strategies.

Full Citation: *Davies, D. R., J. H. Davies, O. Hassan, K. Morgan, and P. Nithiarasu (2007), Investigations into the applicability of adaptive finite element methods to two-dimensional infinite Prandtl number thermal and thermochemical convection, Geochem. Geophys. Geosyst., 8, Q05010, doi:10.1029/2006GC001470.*

Note: *Although this chapter is published in an academic journal, with contributing authors, their contribution was no greater than it would otherwise have been for a regular thesis chapter.*

2.2 Introduction

Over recent decades, numerical modeling has stirred significant interest in the geodynamical community, leading the way in studies of numerous geological processes. This interest is due to the fact that analytical solutions to the various phenomena are normally unavailable, while experimental methods are sometimes time-consuming and often have limitations. The numerical methods that have been employed have generally been based upon finite-difference (e.g. McKenzie et al. 1974, Bodri & Bodri 1978, Matyska & Yuen 2001) and, occasionally, finite-volume techniques (e.g. Tackley 1996, 1998, Ratcliff et al. 1998, Albers

& Christensen 2001), although recently, the finite element method has become more established (e.g. Baumgardner 1985, Farnetani & Richards 1995, Sidorin et al. 1999, Zhong et al. 2000, Zhong 2006). Methods based upon the finite-element approach are attractive since they lead to general-purpose computer codes. However, it is not this feature that has been largely responsible for the recent interest shown in the method, but the fact that it, as indeed is the finite volume method, is based upon an integral formulation, and, hence, is readily implemented on arbitrary discretizations, i.e. unstructured grids. This final point is central to our study.

It is a well known fact that even the use of sophisticated computational models can give inaccurate results if the numerical grid upon which the model is based is unable to capture the significant features of the problem. Indeed, for the large-scale problems encountered in geodynamics, inadequate grid resolution has become a major concern. The majority of phenomena studied (e.g. subduction zone and mid-ocean ridge magmatism) are characterized by the interaction of complex geometries, complex material properties and complex boundary conditions. Such a combination often yields unpredictable and intricate solutions, with narrow regions of high solution gradient frequently found embedded in large areas where the solution varies slowly. It is these high gradient regions that present a serious challenge for computational methods: their location and extent is very difficult to determine *a priori* and, even if their location is identified, with current methods it is often impossible to resolve localized features. The net result is that achieving accurate solutions is a very demanding task for the analyst. Indeed, one of the most challenging problems currently facing geodynamicists is the accurate solution of such ‘multi-scale’ flows.

This issue of accuracy must be balanced with computational considerations. An accurate solution requires that one properly resolves the active features in the simulation. Resolu-

tion, in turn, is related to the number of nodes employed. Obviously, one could generate a solution of high accuracy by employing an extremely fine mesh throughout the computational domain. However, the larger the number of nodes, or degrees of freedom, the greater the demands on computational memory and processing power. Finding the right balance between *accuracy* and computational *efficiency* is therefore a difficult task. Ideally, what is needed is a method capable of yielding an accurate solution, whilst employing as few degrees of freedom as possible.

Standard methods have attempted to achieve this by utilizing non-uniform grids generated *a priori*, with the user exploiting previous experience to define the grid (Davies & Stevenson 1992, Scott 1992). Such methods, however, are not applicable to unsteady problems, since the active regions within the solution domain are constantly mobile and predicting their location at any given time is an impossible task. The majority of previous studies have overcome this issue by employing a uniformly fine grid throughout the computational domain, as described above (e.g. Bunge et al. 2003, Oldham & Davies 2004). This ensures that, as the simulation evolves, active regions are continually in zones of fine resolution and solution accuracy is maintained. Other methods have also been utilized, although more rarely, including both Lagrangian and Arbitrary Lagrangian-Eulerian (ALE) formulations (Fullsack 1995). However, over time, it has become apparent that these methods have their own restrictions and, consequently, a major area of research in the field of geodynamics is the generation of numerical models that accurately portray the nature of the problem, whilst maintaining computational efficiency.

Here, we introduce grid adaptivity, which is a method commonly employed within the field of engineering (Babuska & Rheinboldt 1978, Lohner et al. 1985, Peraire et al. 1987, Pelletier & Ilinca 1995, Nithiarasu & Zienkiewicz 2000). Since we demonstrate the method in the

context of finite elements it is termed the Adaptive Finite Element Method (AFEM). The method provides a powerful approach for the *accurate* and *efficient* solution of the complex problems encountered in geodynamics. Grid points are automatically clustered in regions of rapid solution variation to improve accuracy, leading to a ‘multi-resolution’ solution, with the highest resolutions being analogous to zones of high solution gradient. In simple terms, the method automatically increases or decreases grid resolution where required, leading to more accurate solutions, whilst employing fewer degrees of freedom.

In this study, the method is applied to infinite Prandtl number thermal and thermo-chemical convection in 2-D Cartesian geometry, to investigate its potential benefits within the field of geodynamics. We begin by investigating a benchmark thermal convection problem, with results illustrating that the method is highly successful, improving solution accuracy whilst increasing computational efficiency.

The method is then applied to a more challenging thermo-chemical benchmark problem, employing a grid based method to solve for the compositional field. Our reasons for selecting a grid based method are simple. Recent work (van Keken et al. 1997, Tackley & King 2003) has highlighted the fact that such methods suffer from numerical diffusion, leading to greater entrainment rates in numerical simulations when compared to marker chain and tracer particle methods. This numerical diffusion is predominantly caused by insufficient resolution, a factor naturally addressed by the AFEM. Consequently, two hypotheses are tested:

1. The greater resolution endorsed by the AFEM will reduce artificial diffusion.
2. This reduced diffusion, in turn, will see grid based methods yielding results that are consistent with those achieved using tracer particle and marker chain methods.

As expected, results show that adapted grids yield large improvements over regular, uniform

grids, generating less diffusive results, whilst reducing the number of degrees of freedom. However, results also demonstrate that grid based methods, even when coupled with the AFEM, are not competitive with the other methods currently employed for the tracking of compositional heterogeneities.

The remainder of this chapter will provide an overview of the test problems that are studied, a summary of the numerical techniques employed (an Appendix is included with a more detailed analysis) and a comparison of the results obtained. These demonstrate the applicability of grid adaptivity to the modeling of mantle convection.

2.3 Problem Description

We consider two dimensional thermal and thermo-chemical convection in an isoviscous, infinite Prandtl number Cartesian box. The equations (in dimensionless, vector form) describing such incompressible convection are the equations of:

Momentum:

$$\nabla^2 \mathbf{u} = -\nabla p + RaT\hat{\mathbf{k}} \quad (2.1)$$

Continuity (mass):

$$\nabla \cdot \mathbf{u} = 0 \quad (2.2)$$

and Energy:

$$\frac{\partial T}{\partial t} + \mathbf{u} \cdot \nabla T = \nabla^2 T \quad (2.3)$$

where $\mathbf{u}$ is the dimensionless velocity, T is the dimensionless temperature, p is the dimensionless non-lithostatic pressure, $\hat{\mathbf{k}}$ is the unit vector in the direction of gravity and t is the dimensionless time. In this form, all material properties are combined into one dimensionless

parameter, the Rayleigh number:

$$Ra = \frac{\beta \rho g \Delta T d^3}{\kappa \mu} \quad (2.4)$$

Here, g is the acceleration due to gravity, ρ is density, β is the coefficient of thermal expansion, ΔT is the temperature drop across the domain, d is the domain length, κ is the thermal diffusivity and μ is the dynamic viscosity.

In thermo-chemical simulations, a second advection-diffusion equation is solved for composition:

$$\frac{\partial C}{\partial t} + \mathbf{u} \cdot \nabla C = \frac{1}{Le} \nabla^2 C \quad (2.5)$$

where C represents composition and Le is the Lewis number, the ratio of thermal diffusivity to chemical diffusivity. The desired limit approaches infinite Le , however, for numerical reasons, a finite Le is often assumed. The momentum equation also differs from Equation (2.1) and is now taken in the form:

$$\nabla^2 \mathbf{u} = -\nabla p (RaT - RbC) \hat{k} \quad (2.6)$$

Here, Rb is the compositional Rayleigh number:

$$Rb = \frac{\Delta \rho g d^3}{\kappa \mu} \quad (2.7)$$

where $\Delta \rho$ is the density difference between the dense ($C = 1$) and light ($C = 0$) material.

These equations are solved using a modified version of the widely used 2-D geodynamics finite element program CONMAN. A brief overview of the code is provided here, however, a more detailed description can be found in King et al. (1990).

The momentum and energy equations form a coupled set of differential equations, although the coupling is not strong since the density, ρ , is constant, other than in the buoyancy term

of the momentum equation (Boussinesq approximation). The incompressibility (continuity) equation is treated as a constraint on the momentum equation, with incompressibility enforced using a penalty formulation. The energy equation is solved using a streamline upwind Petrov Galerkin (SUPG) method (Hughes & Brooks 1979), with time stepping accomplished by means of an explicit second order predictor-corrector algorithm.

A grid based method, identical to that used for solving the energy equation, is utilized for simulating the compositional field. A small chemical diffusivity is assumed and a filtering scheme is employed to remove spurious numerical overshoot and undershoot features, which are common with advection diffusion problems of this nature. This filter conserves mass by design and has been shown to work remarkably well at limiting the aforementioned numerical errors (see Lenardic & Kaula 1993, for further details).

2.3.1 Case 1: Thermal Convection in a Square Cavity

The first example considered is buoyancy driven flow in an iso-chemical square cavity. The cavity is filled with a material of constant viscosity and there are no internal heat sources. Boundary conditions are summarized in Figure 2.1. This problem is solved at $Ra = 10^4$, 10^5 , and 10^6 , initially on uniform, structured meshes, and subsequently, on adapted, unstructured meshes.

Data Calculated

The following data or sets of data are calculated during the simulations:

1. The Nusselt Number, i.e. the mean surface temperature gradient:

$$Nu = \int_0^l -\frac{\partial T}{\partial y}(x, y = 1)dx \quad (2.8)$$

where l is the domain length.




Figure 2.1: The boundary conditions utilized when studying thermal convection in a square cavity. T , $V(x)$ and $V(y)$ represent temperature and the velocities in the x and y Cartesian directions respectively.

2. The non-dimensional Root-Mean-Square (RMS) velocity:

$$V_{RMS} = \sqrt{\frac{1}{V} \int_V \|\mathbf{u}\|^2} \quad (2.9)$$

where V is the area of the computational domain.

3. Non-dimensional temperature gradients at domain corners:

$$q = -\frac{\partial T}{\partial y} \quad (2.10)$$

with q_1 at $x = 0$, $y = 1$; and q_2 at $x = y = 1$.

2.3.2 Case 2: Entrainment of a Thin Dense Layer through Thermo-chemical Convection

The second example considered is a well-established thermo-chemical benchmark problem from van Keken et al. (1997). We model the entrainment of a deep-seated, thin (0.025), dense layer, in an aspect ratio 2 box, of unit height. This layer is prescribed a compositional value $C = 1$, while the overlying, lighter material has a value of $C = 0$. The problem is analogous to the entrainment of a compositionally dense layer in the D'' region at the Core-Mantle boundary. Results are computed for $1/Le = 10^{-6}$. The thermal Rayleigh Number, Ra , is set to 3×10^5 , the compositional Rayleigh Number, Rb , is set to 4.5×10^5 , while the




Figure 2.2: The boundary conditions utilized when studying the entrainment of a thin dense layer, in an aspect ratio 2 box, through thermo-chemical simulations. T , C , $V(x)$ and $V(y)$ represent temperature, composition and the velocities in the x and y directions respectively.

viscosity is assumed to be constant. Boundary conditions are summarized in Figure 2.2. As an initial condition, an analytical expression of the temperature based on boundary layer theory is taken (see Appendix A in van Keken et al. 1997). Once again, this problem is solved initially on uniform, structured meshes, and subsequently, on adapted, unstructured meshes.

Data Calculated

For this case, the relative entrainment (e) is calculated as a function of time, from:

$$e = \frac{1}{\lambda d_b} \int_{d_e}^l C dV \quad (2.11)$$

where λ is the aspect ratio of the box ($= 2$), d_b is the thickness of the dense layer and d_e is an arbitrarily chosen height ($= 0.2$ here). We focus our attention on the relative entrainment as opposed to other parameters, such as RMS velocity, since the rate of entrainment provides an excellent means to track the evolution of a thermo-chemical simulation. Additionally, entrainment rates have been calculated in previous studies (van Keken et al. 1997, Tackley & King 2003) and, hence, direct comparisons can be made with ease.

Before providing a summary of the adaptive strategies employed, we should point out that the two cases presented above not only provide a means to test the applicability of the AFEM to thermal and thermo-chemical convection problems, but they also allow us to test the AFEM in both steady state (Case 1) and unsteady (Case 2) situations. Both are common within the field of geodynamics.

2.4 Adaptive Methodology

As is clear from the introduction, the methods employed in modeling thermal and thermo-chemical convection must provide an adequate definition of the problem, in a computationally efficient manner. In other words, the methods must be adept at resolving narrow regions of high gradients that frequently occur and are normally found embedded in large areas where the flow variables vary slowly. Since the exact location of these high gradient regions are not always known to the analyst *a priori*, particularly with unsteady problems, it is apparent that adaptive mesh methods, with *a posteriori* error indicators, could have an important role to play in the development of efficient solution techniques for such problems. At present, a wide variety of adaptive procedures are being utilized within the engineering community. Broadly speaking, these fall into two categories:

1. *h*-refinement, in which the same class of elements continue to be used but are changed in size, in some locations made larger, and in others made smaller, to provide the maximum economy in reaching the desired solution.
2. *p*-refinement, in which the same element size is utilized, but the order of the polynomial is increased or decreased as required (e.g. linear shape functions are ‘adapted’ to quadratic or higher order).

A variant of the *h*- method, known as ‘adaptive remeshing’, is employed in this study. It provides the greatest control of mesh size and grading to better resolve the flow features.

In this method, for steady state simulations, the problem is solved initially on a grid fine enough to roughly capture the physics of the flow. Remeshing then involves the following steps:

1. The solution is analyzed through some kind of error indication procedure, to determine locations where the mesh fails to provide an adequate definition of the problem. This can either mean that additional grid points are needed, or perhaps, that there are too many grid points at certain locations within the domain. An interpolation based local error indicator is employed in this study, based upon nodal solution gradients and curvatures.
2. The information yielded by this error estimation process is utilized to generate an improved mesh through an automatic mesh generator. A variant of the so-called advancing front technique is utilized here, being capable of generating meshes that conform to a user prescribed spatial distribution of element size. Elements can locally increase or decrease in size as required, leading to what we term an ‘optimal’ mesh.
3. The original solution is interpolated between meshes using higher-order cubic interpolation (Nielson 1979, El Hachemi et al. 2003).
4. The solution procedure continues on the new mesh.

The remeshing process is repeated until a desired level of accuracy has been achieved.




Figure 2.3: A flow diagram summarizing the major steps involved in the adaptive remeshing procedure.

For unsteady problems, the process is almost identical, however, it is fundamental that the initial mesh is suitably defined. If the mesh is inadequate, the errors generated during the calculation's early stages propagate through the computational domain, generating misleading results. To ensure such errors do not arise in our simulations, we generate an optimal initial mesh. This process is straightforward. The initial condition is set up on a structured grid. In the same way as is described in points 1-4 above, the data from this structured grid is analyzed to determine where the mesh needs modification. This information is then used to regenerate the mesh, and the solution (i.e. the initial condition) is transferred onto the new mesh using cubic interpolation. Consequently, the initial grid naturally provides an optimal definition of the problem. With unsteady problems, the remeshing procedure can be continued indefinitely as the simulation evolves. The remeshing 'loop' is activated after a user defined time interval, or dynamically, based upon an approximation to the error. It is important to point out that although the process seems similar to Lagrangian formulations (the computational mesh appears to follow the solution), it is indeed an Eulerian formulation, perhaps with some of the advantages commonly associated with Lagrangian schemes.

For a more detailed description of the mesh generation process, the error estimation procedure and the adaptive strategy, please refer to Appendices C and D. A flow chart summarizing the essential stages involved is also included in Figure 2.3.

2.4.1 Remeshing Procedure for Case 1 - Steady-State Adaptivity

A steady-state solution is achieved here. Remeshing is therefore a simple task and is performed when the solution converges to a steady state on a given grid. The process is terminated when an optimal mesh has been produced, i.e. the solution does not improve with the remeshing procedure. The error indicator employed in this case is based upon

	Mesh 1	Mesh 2	Mesh 3	Mesh 4	Mesh 5	Mesh 6	Benchmark	Uncertainty
Elements	256	1024	4096	7225	10000	16384	****	****
Nu	7.8757	9.6838	10.3361	10.4453	10.4712	10.5058	10.5341	0.00001
V_{RMS}	197.4998	194.2366	193.4223	193.3222	193.2968	193.2787	193.2145	0.0001
q_1	11.8004	16.5079	18.3362	18.6467	18.7641	18.8861	19.0794	0.00004
q_2	1.4204	0.9397	0.7754	0.7529	0.7445	0.7358	0.7228	0.00002

Table 2.1: The results obtained on various uniform, structured meshes for simulations of thermal convection in a square cavity at $Ra = 10^5$. The benchmark results of Blankenbach et al. (1989) are included for comparison, together with their uncertainties.

nodal temperature gradients and curvatures.

2.4.2 Remeshing Procedure for Case 2 - Temporal Adaptivity

The results presented are based upon simulations with a remeshing frequency of 2000 time steps. This value was selected after a series of tests, both visual and analytical, tracing the temporal evolution of the model. Ideally, the remeshing procedure would be linked to the dynamics (i.e. it should be tied to some measure of how much the solution has changed or whether derivatives in the mesh exceed a certain tolerance). Nonetheless, we have verified, through tests at various remeshing frequencies, that for this simulation, the remeshing frequency selected (i.e. 2000 time steps) does not degenerate the results.

The error indicator employed is similar to that in Case 1, however it is based upon a combination of temperature and composition, as opposed to temperature alone. Nodal solution derivatives are calculated for both the thermal and compositional fields. The highest values yielded are then selected as derivatives for that particular node. Such a scheme engenders high resolving power at the density interface, as well as sufficient resolution to accurately solve the thermal field. We have verified that this combination yields superior results to simulations employing a combination of the composition and velocity variables.

The remeshing strategy is slightly different to that of Case 1. Rather than simply refining zones of high solution gradient, we also allocate fine resolution to neighboring regions.

	UM	Benchmark	Uncertainty
Elements	16384	****	****
Nu	4.8952	4.8844	0.00001
V_{RMS}	42.8713	42.8649	0.00002
q_1	8.0457	8.0594	0.000003
q_2	0.5905	0.5888	0.000003

Table 2.2: The results obtained on uniform meshes (UM) for simulations of thermal convection in a square cavity at $Ra = 10^4$.

This is done by ensuring that the transition from fine to coarse elements is extremely gradational. In the steady-state cases, we specify that a minimum of 10 elements is required to make this transition. However, in our thermo-chemical, unsteady simulations, this value is set to 30. In this way, fine regions of the mesh are always surrounded by relatively fine zones. Accordingly, as the simulation evolves, accuracy is maintained, since zones of high solution gradient have not departed the fine grid regions before the next remeshing. Such a scheme allows a greater time interval between remeshings and was a key consideration when selecting the remeshing frequency.

It is important to note that the minimum element size (δ_{min}) permitted in our simulations is ≈ 0.002 which yields localized resolution equivalent to that achieved during a uniform mesh simulation of 1000×500 elements, in a box of aspect ratio 2.

2.5 Results & Discussion

2.5.1 Thermal Convection in a Square Cavity

Uniform Structured Meshes

The results obtained for uniform mesh calculations at $Ra = 10^5$ are displayed in Table 2.1. The results at both $Ra = 10^4$ (Table 2.2) and $Ra = 10^6$ (Table 2.3) demonstrate a comparable relationship, and consequently only one set of results are presented fully.

	UM	Benchmark	Uncertainty
Elements	16384	****	****
Nu	21.3773	21.9725	0.00002
V_{RMS}	834.9486	833.9898	0.0002
q_1	43.3217	45.9673	0.0003
q_2	0.9737	0.8772	0.00001

Table 2.3: The results obtained on uniform meshes (UM) for simulations of thermal convection in a square cavity at $Ra = 10^6$.

Rayleigh Number	Solution Error (%)
10^4	0.2
10^5	0.9
10^6	4.9

Table 2.4: The solution errors obtained on a uniform mesh of 16384 elements (128×128 elements), at various Rayleigh Numbers, for thermal convection in a square cavity.




Figure 2.4: The relationship between the number of elements in a mesh and (a) RMS velocity / (b) mean Nusselt Number. The results represent thermal convection in a square cavity, on uniform, structured meshes, at $Ra = 10^5$. Benchmark values are represented by horizontal dashed lines.

Figure 2.4 displays the relationship between the number of elements in a mesh and the RMS velocity and mean Nusselt Number respectively. The figure illustrates that a large number of elements (i.e. > 4000) is required before results begin to show some sort of consistency. Even more elements are required before results converge on the benchmark solutions of Blankenbach et al. (1989).

By calculating the percentage error for each output analyzed (i.e. mean Nu , RMS velocity, q_1 and q_2) and subsequently taking the mean of these four percentages, we have determined the discrepancy, i.e. the solution error, between uniform mesh solutions and benchmark solutions at various Rayleigh numbers. Results are summarized in Table 2.4.

As convection intensifies, the solution error becomes more prominent, as would be expected. At higher Rayleigh numbers (10^5 , 10^6), RMS velocities, mean Nusselt numbers, and corner temperature gradients lie far beyond the realms of uncertainty of the benchmark solution. Even at $Ra = 10^4$, solutions fail to achieve a satisfactory level of accuracy. However, as we show next, with the use of adaptive, unstructured meshes, solution accuracy is greatly enhanced.

Adapted Meshes

The results obtained at $Ra = 10^4$, 10^5 and 10^6 alongside their final adapted meshes are displayed in Figure 2.5. They are also summarized in Tables 2.5 – 2.7. For completeness, the sequence of meshes employed at $Ra = 10^5$ are displayed in Figure 2.6, together with the corresponding temperature contours. Once again, results for both $Ra = 10^4$ and 10^6 display a comparable relationship and, consequently, only one set of results is presented fully.




Figure 2.5: Final adapted meshes and corresponding temperature contours for purely thermal convection in a square cavity at $Ra = 10^4$, 10^5 and 10^6 . Red is hot ($T=1$), blue is cold ($T=0$), and the contour spacing is 0.02.

The characteristics of each mesh at $Ra = 10^5$ are summarized in Table 2.8, which shows the number of quadrilateral elements and nodal points, and the values of the generation parameters (δ_{Min} , δ_{Max} , s_{Max} and C) defined in Appendix C. It is important to note that for consistency, the generation parameters displayed in Table 2.8 are also utilized in our simulations at $Ra = 10^4$ and $Ra = 10^6$. It should also be noted that we have intentionally restricted the number of elements to less than 16,500, to allow for a simple comparison between adapted, unstructured meshes, and uniform, structured meshes.

As seen, the proposed procedure has refined locations of thermal boundary layers, wherever they are strong. Consequently, when compared to uniform structured meshes, the majority of results show far superior concurrence with the benchmark solution, in spite of a reduction in the number of elements (Table 2.9). As noted previously, at $Ra = 10^5$, a uniform mesh of more than 16,000 elements is required for a reasonably well-resolved solution, i.e. within 1% of the benchmark results. At $Ra = 10^6$, extrapolating from a series of uniform mesh simulations, we expect that more than 50,000 elements would be required before results converge to within $\approx 4\%$ of the benchmark solution. More accurate results are inaccessible with the lower-order, serial configuration of CONMAN. However, with the proposed adaptive procedure this is not the case. At $Ra = 10^5$, solutions within 0.1% of the benchmark are achieved on a mesh of $\approx 15,700$ elements and at $Ra = 10^6$, results converge to within 1% on a mesh of $\approx 16,200$ elements. Using the AFEM, the number of elements required for adequate solution varies with Rayleigh number, but is always significantly less than that of a uniform mesh for a specified precision.

It is important to point out that the benefits of this technique become more noticeable when convection is more vigorous and temperature gradients are greater. At $Ra = 10^4$, only a moderate increase in accuracy is observed between adapted and uniform meshes, for

	Mesh 1	Mesh 2	Mesh 3	Mesh 4	Benchmark
Elements	412	1977	6474	14972	****
Nu	4.2687	4.8565	4.8683	4.8790	4.8844
V_{RMS}	40.8536	42.9019	42.8792	42.8679	42.8649
q_1	7.2917	8.0345	8.0479	8.0538	8.0594
q_2	0.6366	0.5887	0.5886	0.5885	0.5888

Table 2.5: The results obtained after each remeshing for simulations at $Ra = 10^4$ on non-uniform, adapted meshes.

	Mesh 1	Mesh 2	Mesh 3	Mesh 4	Benchmark
Elements	412	2390	7321	15722	****
Nu	7.9675	10.0967	10.4380	10.5278	10.5341
V_{RMS}	184.8429	193.5598	193.3056	193.2456	193.2145
q_1	13.1450	18.9078	19.0353	19.0518	19.0794
q_2	1.2687	0.7301	0.7195	0.7216	0.7228

Table 2.6: The results obtained after each remeshing for simulations at $Ra = 10^5$ on non-uniform, adapted meshes.

	Mesh 1	Mesh 2	Mesh 3	Mesh 4	Benchmark
Elements	412	2733	7681	16195	****
Nu	11.6429	18.5597	20.9920	21.5077	21.9725
V_{RMS}	809.9980	841.9425	835.1100	834.7003	833.9898
q_1	12.3555	38.3885	44.7377	45.601	45.9643
q_2	2.6031	0.9379	0.9038	0.8694	0.8772


Table 2.7: The results obtained after each remeshing for simulations at $Ra = 10^6$ on non-uniform, adapted meshes.

Mesh	Elements	Nodes	δ_{Min}	δ_{Max}	s_{Max}	C
1	412	453	0.05	0.05	1.0	-
2	2390	2493	0.009	0.05	5.0	0.6
3	7321	7469	0.0045	0.030	5.0	0.3
4	15722	15961	0.0025	0.025	5.0	0.2


Table 2.8: The sequence of meshes employed for the problem of buoyancy driven flow in a square cavity at $Ra = 10^5$.

Rayleigh Number	Elements	Solution Error (%)
10^4	14972	0.06
10^5	15722	0.09
10^6	16195	1




Table 2.9: The solution errors yielded by non-uniform adapted meshes at various Rayleigh Numbers.




Mesh 1 – 453 Nodes, 412 Elements.



Mesh 2 – 2493 Nodes, 2390 Elements.



Mesh 3 – 7469 Nodes, 7321 Elements.



Mesh 4 – 15961 Nodes, 15722 Elements.





Figure 2.6: The series of meshes employed at $Ra = 10^5$ along with corresponding temperature contours. Red is hot ($T=1$), blue is cold ($T=0$), and the contour spacing is 0.02.

approximately the same number of elements (solution error decreases by a factor of 3, from $\approx 0.2\%$ to $\approx 0.06\%$). However, at $Ra = 10^6$, results display a superior concordance with the benchmark solution on adapted meshes (the solution error decreases by a factor of 5, from $\approx 5\%$ to $\approx 1\%$). Indeed, on fully uniform meshes with linear shape functions, obtaining an accurate solution on a single processor would be highly impractical at $Ra = 10^6$.

Other conclusions can be drawn by analyzing the 4 outputs individually. The remeshing process clearly has a positive effect on the global measures (i.e. V_{RMS} and Nu). However, the procedure seems to have a more dramatic effect on the accuracy of the heat flux at domain corners (q_1 and q_2), particularly at higher Rayleigh numbers. This is easy to understand - the heat flux at these corners is strongly influenced by the resolution achieved in the upper thermal boundary layer. Since this boundary layer is characterized by high temperature gradients, the remeshing procedure refines the grid significantly in these regions. Accordingly, the corner solutions yielded by adapted meshes are far superior to those yielded on uniform structured meshes. Global measures on the other hand, are not influenced by these boundary layers to such an extent. Consequently, the improvement observed in global measures between adapted and uniform grids is less dramatic.

Processing Efficiency

We have demonstrated that the number of nodes and elements required for accuracy is less with the AFEM. Consequently, the AFEM is more efficient in terms of memory requirements (RAM). However, is the AFEM economical in terms of CPU-time?

Figure 2.7 illustrates the relationship between solution error and the time taken in obtaining these solutions on both uniform and adapted meshes (at various Rayleigh numbers). It should be noted that the timings displayed for the adaptive cases include the time allocated

Rayleigh Number	% Time Remeshing
10^4	4
10^5	1.7
10^6	0.8

Table 2.10: The percentage of computational time taken by the remeshing process at various Rayleigh Numbers.

for remeshing. The main points arising are summarized below:

1. In general, at $Ra = 10^4$, the AFEM is less efficient than uniform meshes for a prescribed level of accuracy. However, when a solution error of less than $\approx 0.2\%$ is required, the AFEM becomes more economical.
2. At $Ra = 10^5$ and $Ra = 10^6$ the AFEM is more efficient than uniform structured meshes, decreasing computational processing time whilst increasing solution accuracy. Indeed, the lower graph in Figure 2.7 ($Ra = 10^6$) suggests that solution errors of less than $\approx 4\%$ cannot be obtained on uniform meshes. With the AFEM this is not the case - far superior results are attained, although due to the lower order nature of CONMAN, minor errors persist.

Table 2.10 summarizes the percentage of calculation time taken by the remeshing procedure, compared to the time spent solving the governing equations. It is important to note that although remeshing appears more efficient at higher Rayleigh numbers, this is not strictly true. At higher convective vigors, simulations take longer to converge towards a steady state solution. Since the number of remeshing loops employed during each simulation is constant (i.e. 3), it is apparent that remeshing will take a smaller fraction of the calculation time at higher Rayleigh numbers (since the calculation time as a whole is greater). Regardless of this point, the main conclusion to be drawn from the data is that the remeshing procedure is computationally inexpensive, expending only a small percentage of the calculation time.

In summary, in the context of purely thermal convection, the number of degrees of freedom required for accuracy on uniform structured meshes is greater than that required for adapted




Figure 2.7: The time taken to converge on various solution errors with both uniform (continuous line) and adapted (dashed line) meshes at (a) $Ra = 10^4$, (b) $Ra = 10^5$ and (c) $Ra = 10^6$.

meshes. Thus, for the same precision, the number of nodes and elements is reduced when the adaptive procedure is used. Additionally, the remeshing procedure is computationally inexpensive and, consequently, particularly at higher Ra ($> 10^5$), the AFEM allows one to attain a desired solution in less processing time.

2.5.2 Thermo-chemical Convection

Having demonstrated the applicability and benefits of the AFEM for thermal convection, we move on to thermo-chemical simulations.

Uniform Structured Meshes

Curves illustrating the entrainment (e) yielded by uniform mesh simulations at various grid resolutions are displayed in Figure 2.8. The results are consistent with previous work in that relative entrainment decreases with increased resolution. As a quantitative example, entrainment at $t = 0.02$ decreases from a value of 0.4 on a grid of 64×64 elements, to a value of 0.08 on a grid of 256×256 elements. This large reduction can be easily understood - grid based methods suffer from numerical diffusion, which smears density interfaces over several grid spacings. Consequently, with coarser grids, material at the density interface itself becomes easier to entrain. As one increases grid resolution, grid spacings decrease, leading to less artificial diffusion and, hence, reduced entrainment.

Perhaps a more fundamental point to note from Figure 2.8, however, is that even after this significant increase in resolution, when $t = 0.01$, and indeed for a high percentage of the calculation, the grid based method displays almost an order of magnitude greater entrainment than the other methods currently employed in this field. The marker chain results produced by Christensen, Neumeister and Dion (CND) from van Keken et al. (1997) are displayed for comparison. The reader should note that the tracer particle studies of




Figure 2.8: (a) Relative entrainment (e) against time on a series of uniform meshes. The 'best' results from van Keken et al. (1997) - CND Markerchain - are also displayed for ease of comparison. (b) An enlargement of the results in (a) for time ≤ 0.02 .




Figure 2.9: Seven figures, at regular time intervals of 0.01, illustrating the evolution of the compositional field, modeled on a uniform mesh of 256×256 elements. Red represents dense material ($C = 1$), while blue represents lighter material ($C = 0$).

Tackley & King (2003) are consistent with these results. Clearly, there is no resemblance between the grid based methods and these results, particularly during the early stages of the simulation. At first glance, it does appear that when $t > 0.02$, the results yielded by grid based simulations are reasonably consistent with those of previous particle studies. Closer analysis however, reveals that this is misleading. The drastic difference in entrainment during the early development of these models means that, by this time, both simulations have evolved into completely different problems. Consequently, the apparently consistent relationship between both methods is purely coincidental. This point is reinforced when studying the visual patterns - there is only a poor resemblance between Figure 2.9 here, which shows the temporal evolution of the compositional field, and equivalent Figures 8 & 10 in van Keken et al. (1997) and Figure 3 in Tackley & King (2003). The similarities also diminish as the simulation evolves. The main differences include the position of the dense pile after reorganization into a two cell pattern and the amount of material trapped at the stagnation point. An important observation to make is that, during the later stages of the calculation, diffusion dominates and, by $t = 0.07$, the dense layer has virtually disappeared. This is not the case when tracer particle and marker chain methods are employed, as has been pointed out by van Keken et al. (1997).

It is clear from the simple tests performed here, on uniform meshes, that increased grid resolution reduces entrainment. Results move towards those yielded with tracer particle and marker chain methods, however, entrainment rates remain significantly higher. Next, we will attempt to answer whether this is the case when the AFEM is employed or does the method provide sufficient resolution to resolve these discrepancies?

δ_{Min}	δ_{Max}	s_{Max}	C
0.002	0.02	5	0.1

Table 2.11: Mesh parameters employed for thermo-chemical entrainment simulations.

Adapted Unstructured Meshes

The generation parameters (δ_{Min} , δ_{Max} , s_{Max} and C) utilized within these models are summarized in Table 2.11. Unlike those previously defined for the purely thermal case (Table 2.8), the values remain constant throughout the simulation. There is a simple explanation for this. Since the problem does not converge to a steady-state solution, one must ensure that not only is the error equally distributed spatially, but also temporally. By fixing the generation parameters over time, one guarantees that this is the case.

Entrainment curves from these simulations are displayed in Figure 2.10. The result obtained from a uniform mesh simulation on a grid of 256×256 elements is also displayed for ease of comparison, along with the ‘best’ results (CND) of van Keken et al. (1997).

Plots are displayed from simulations employing linear (AM-Linear) and cubic (AM-Cubic) interpolation between grids. The accuracy of this interpolation process is of fundamental importance during unsteady problems of this nature. With steady state problems, the solution always converges towards a certain end member. Consequently, minor errors arising due to inaccurate interpolation can be overcome. However with unsteady problems, errors arising during the interpolation process propagate through the computational domain, leading to a solution that is unrepresentative of the true problem. In essence, the solution emerging from the remeshing procedure must be exactly that which enters. Otherwise, the simulation evolves falsely.

The results involving linear interpolation are, therefore, only presented for completeness




Figure 2.10: (a) Relative entrainment (e) against time on adapted, unstructured meshes. The results obtained on a 256×256 element uniform mesh are also displayed, as well as the ‘best’ results from van Keken et al. (1997). (b) An enlargement of the results in (a) for time ≤ 0.02 .




Figure 2.11: Evolution of the compositional field (left) and temperature field (right) on a series of adapted grids (center). The grids are adapted around temperature and compositional solution gradients. Additionally a region of fine resolution is generated adjacent to zones of high solution gradient. Consequently, as the simulation evolves, high gradient zones remain in regions of high resolution, leading to less numerical error.

- the method fails to accurately interpolate features at the density interface. At certain locations, the density jumps from a value of 0 to 1 within a single element. Linear interpolation is not capable of resolving such a feature and, consequently, the remeshing process employing linear interpolation generates diffusive and non-conservative results. The cubic interpolation strategy employed, however (Nielson 1979, El Hachemi et al. 2003), accurately captures such features, being both locally and globally conservative.

The results demonstrate that in general, the AFEM leads to a significant reduction in artificial diffusion and, hence, entrainment rates, when compared to uniform mesh simulations. This point is reinforced by making a comparison of the longevity of the dense pile in the visual output (Figures 2.9 & 2.11). As was noted earlier, with a uniform mesh simulation on a grid of 256×256 elements, by $t = 0.07$, the dense pile has almost completely diffused. This is not the case with simulations employing the AFEM – the dense pile remains extremely coherent until this time. It is clear, therefore, that the higher resolution permitted by the AFEM leads to decreased artificial diffusion and, consequently, results that provide a more precise representation of the problem.

It is important to point out that, in addition to yielding superior results, our adaptive grid simulations are computationally more efficient, both in terms of memory and processing time, when compared to uniform mesh simulations. The number of elements utilized in this simulation varies with time between $\approx 40,000$ and $\approx 58,000$, depending upon the configuration of the calculation. Additionally, the generation of a new optimal mesh is an inexpensive procedure, typically taking between 15 and 20 time-steps, compared to the time taken for one time step with a non-changing mesh (although the time expended in remeshing can be decreased significantly by specifying a larger *err_max* - the remeshing tolerance - see Appendix C, section C.2.2). Obviously, this gain in computational efficiency is only valid

provided one does not set the minimum element size, δ_{min} , to an unreasonably small value. Our experience suggests that simulations with $\delta_{min} < 0.002$ are highly impractical, due to the tradeoff between minimum element size and time stepping, dictated by the Courant-Friedrichs-Levy condition. Of course, this situation could be remedied by employing a local time stepping algorithm, however, this was beyond the scope of our study. It should be noted that the efficiency of the method as a whole depends on how often it is necessary to remesh, which depends on how time-dependent the simulation is. Mesh adaptivity therefore becomes less efficient for highly time-dependent cases needing frequent remeshing.

Although the results presented demonstrate the benefits of the AFEM, grid-based thermo-chemical methods, even when coupled with the AFEM, yield results that are beyond the realms of uncertainty of those achieved via tracer particle and marker chain methods. The method remains particularly diffusive, yielding erroneous entrainment rates, even at the extremely fine resolutions permitted in our simulations. Once again, these findings are reinforced by the visual patterns, with Figure 2.11 displaying only a marginal resemblance to Figures 8 & 10 in van Keken et al. (1997) and Figure 3 in Tackley & King (2003). This resemblance however, is stronger than that observed with our uniform mesh simulations (Figure 2.9).

2.6 Conclusions

An adaptive finite element procedure has been presented for solving convective heat transfer problems within the field of geodynamics. The method adapts the mesh automatically around regions of high solution gradient, yielding enhanced resolution of the associated flow features.

2.6.1 Applicability to Thermal Convection

The results obtained from thermal convection simulations are extremely positive. The error indicator presented has proven reliable and the adaptive procedure is shown to be robust. Predictions for heat transfer agree well with benchmark solutions, suggesting that the technique is valid and accurate.

2.6.2 Applicability to Thermo-chemical Convection

The results obtained from thermo-chemical simulations are somewhat less conclusive. However, two major conclusions can be drawn:

1. The AFEM provides a suitable means for increasing grid resolution in localized regions. This leads to a reduction in numerical diffusion and, hence, entrainment rates, provided that the interpolation employed during the remeshing procedure accurately captures all underlying features. Our results suggest that an extension of this work to both tracer particle and marker chain methods would be a worthwhile exercise, with the higher spatial resolution yielded leading to the more accurate tracking of particles (or the marker chain), generating results of greater accuracy.
2. Even using the AFEM, grid based methods fail to achieve results that are consistent with other methods. Consequently, we conclude that the method, at least in its current format, requires unrealistically high resolution to limit artificial diffusion and accurately track chemical heterogeneities.

In summary, the number of degrees of freedom required for accuracy on uniform structured meshes is greater than that on adapted unstructured meshes. Thus for the same, or often superior precision, the number of degrees of freedom is reduced when the present adaptive procedure is used. However, perhaps the most important advantages of the AFEM are that:

1. The unstructured nature of the technique allows its use when modeling many of the complex geometries encountered on Earth.

2. Nodes automatically cluster around zones of high solution gradient, without the need for complicated *a priori* mesh generation.
3. The reduction in the number of degrees of freedom leads to a decrease in CPU-time and memory use (both in terms of disk space and RAM), meaning that complex problems can be solved efficiently.

To date, successful goal-orientated/error-guided grid adaptation techniques have, to our knowledge, not been utilized in the field of geodynamics. Potential applications of the method are wide-ranging and specific elements of the method could even be applied alone in certain situations. For example, the error-guided remeshing procedure would prove extremely useful in Lagrangian simulations, when large distortions of the computational domain necessitate a total regridding. Within the field of geodynamics applications of the AFEM would include studies into subduction zone dynamics, D'' and its interaction with the post-perovskite phase transition, upper mantle phase transitions, mid-ocean ridge magmatism and plume dynamics. These phenomena have one thing in common - 'active' regions of high solution gradient, be it in temperature, composition, pressure or velocity, that are found embedded in large, 'passive' regions, whose location is difficult to determine *a priori*. It is clear therefore, that adaptive grid methods, with *a posteriori* error indicators, should have an important role to play in the development of efficient solution techniques for such problems. This development should not be restricted to the method described here (i.e. adaptive remeshing), but to adaptive procedures as a whole. Due to its flexibility, numerical modeling will undoubtedly continue as a primary tool in helping us to understand various geological processes. The AFEM and the 'multi-resolution' solutions yielded, should, for the moment at least, ensure that progress is not unnecessarily restricted by computer power.

Chapter 3

Adaptive Finite Element Methods in Geodynamics; Convection Dominated Mid–Ocean Ridge and Subduction Zone Simulations

3.1 Abstract

An adaptive finite element procedure is presented for improving the quality of convection dominated mid–ocean ridge and subduction zone simulations in geodynamics. The method adapts the mesh automatically around regions of high solution gradient, yielding enhanced resolution of the associated flow features. The approach utilizes an automatic, unstructured mesh generator and a finite element flow solver. Mesh adaptation is accomplished through mesh regeneration, employing information provided by an interpolation based local error indicator, obtained from the computed solution on an existing mesh. The proposed method-

ology works remarkably well at improving solution accuracy for both mid-ocean ridge and subduction zone simulations. Furthermore, the method is computationally highly efficient. To date, successful goal-orientated/error-guided grid adaptation techniques have, to our knowledge, not been utilized within the field of geodynamics. This chapter presents the first true geodynamical application of such methods.

Full Citation: *Davies, D. R., J. H. Davies, O. Hassan, K. Morgan, and P. Nithiarasu (2007), Adaptive Finite Element Methods in Geodynamics; Convection Dominated Mid-Ocean Ridge and Subduction Zone Simulations, Int. J. Num. Meth. Heat Fluid Flow., In Press.*

Note: *Although this chapter is published in an academic journal, with contributing authors, their contribution was no greater than it would otherwise have been for a regular thesis chapter.*

3.2 Introduction

Over recent decades, adaptive grid techniques (Babuska & Rheinboldt 1978, Lohner et al. 1985, Peraire et al. 1987) have been widely employed by the engineering community, in areas ranging from compressible aerodynamics (e.g. Hassan et al. 1995) to incompressible flow and heat transfer problems (e.g. Pelletier & Ilinca 1995, Nithiarasu & Zienkiewicz 2000, Mayne et al. 2000). However, until recently (Davies et al. 2007), grid adaptivity had not been applied within the field of geodynamics, a branch of geophysics concerned with measuring, modeling, and interpreting the configuration and motion of Earth's crust and mantle. This is surprising, since the method provides a suitable means to solve many of the complex problems currently encountered in the field. The motivation behind this study, therefore, is to demonstrate the benefits of such techniques within a geophysical framework.

The mantle, the region between Earth's crust and core, contains 84% of Earth's volume

and 68% of its mass, but because it is separated from direct observation by the thin crust there are many unsolved problems. Mantle convection establishes one of the longest time scales of our planet. Earth’s mantle, though solid, is deforming slowly by a process of viscous creep and, while sluggish in human terms, the rate of this subsolidus convection is remarkable by any standard. Indeed, it is estimated that the mantle’s Rayleigh number, a dimensionless parameter quantifying its convective instability, is of order 10^9 (Schubert et al. 2001), generating flow velocities of $2\text{--}10\text{cm yr}^{-1}$. Plate tectonics is the prime surface expression of this convection (e.g. Davies & Richards 1992), although, ultimately, all large scale geological activity and dynamics of the planet, such as mountain building and continental drift, involve the release of potential energy within the mantle. Consequently, innovative techniques for simulating these large scale, infinite Prandtl number convective systems are of great importance.

Rather than simulate the whole mantle, which would require massively parallel codes in 3-D spherical geometry, this investigation focusses upon geologically active regions along Earth’s surface, where the mantle interacts with Earth’s crust. Steady state thermal convection is examined, at a Mid–Ocean Ridge (MOR) and at a Subduction Zone (SZ), problems that can be well approximated in two dimensions. A MOR is a long, elevated volcanic structure, occurring at divergent plate margins along the middle of the ocean floor. Such ridges form through the symmetrical spreading of two tectonic plates from the ridge axis. SZ, on the other hand, occur at convergent plate margins, where Earth’s tectonic plates move towards each other, with one plate subducting beneath the other into Earth’s mantle. The geometry of a SZ is mapped out by the locations of earthquakes and deep seismicity, with most present day SZs extending from trenches on the ocean floor, at an angle ranging from near horizontal to near vertical, to a depth of up to 700 km. Volcanoes tend to form $\approx 100\text{km}$ above the subducting slab, at the volcanic arc, making SZs the most active tectonic locations on our

planet.

Numerical simulations of these tectonic settings involve complex geometries, complex material properties and complex boundary conditions. Such a combination often yields unpredictable and intricate solutions, where narrow regions of high solution gradient are found embedded in a more passive background flow. These high gradient regions present a serious challenge for computational methods: their location and extent is difficult to determine *a priori*, since they are not necessarily restricted to the boundary layers of the domain. Furthermore, even if their location is identified, with the current methods employed in the field, it is often impossible to resolve localized features. It is natural to think, therefore, that grid adaptivity, with *a posteriori* error indication criterion, could play an important role in the development of efficient solution techniques for such problems.

The present study extends on the work of Chapter 2, which applied grid adaptivity to infinite Prandtl number, thermal and thermo-chemical convection. However, here, attention is focussed on geodynamical application, as opposed to methodology formulation and validation. The aim is to improve the solution quality of MOR and SZ simulations, by utilizing adaptive mesh refinement strategies. Results illustrate that the method is advantageous, improving solution accuracy whilst reducing computational cost.

The remainder of this chapter will cover the equations governing mantle convection, together with the numerical and adaptive strategies used in their solution. The methodology is then applied in geodynamical simulations of the tectonic settings introduced above.

3.3 Methodology

3.3.1 Governing Equations and Solution Procedure

Earth's mantle is solid. However, over large timescales it deforms slowly through processes such as dislocation and diffusion creep. As a consequence, motion within Earth's mantle can be described by the equations governing fluid dynamics. Since the mantle has an extremely large viscosity ($\approx 10^{21}$ Pa s), the equations governing mantle convection are somewhat different to those governing the more typical fluid mechanics problems:

1. The large viscosity of Earth's mantle makes the Prandtl Number (Pr), the ratio between viscous and inertial forces, of the order 10^{24} . Accordingly, inertial terms in the momentum equation can be ignored.
2. The Ekman number (i.e. the ratio between viscous and Coriolis forces) is of the order 10^9 , since the velocity of convection within the mantle is so small. As a consequence, the Coriolis force can be neglected.
3. The centrifugal force is proportional to the square of the velocity. Consequently, it is even smaller than the Coriolis force and it is also ignored.

This mantle convection problem is formulated in terms of the conservation equations of momentum, mass and energy, expressed for incompressible, Boussinesq convection, in dimensionless, vector form:

$$\nabla^2 \mathbf{u} = -\nabla p + Ra T \hat{\mathbf{k}} \quad (3.1)$$

$$\nabla \cdot \mathbf{u} = 0 \quad (3.2)$$

$$\frac{\partial T}{\partial t} + \mathbf{u} \cdot \nabla T = \nabla^2 T \quad (3.3)$$

where $\mathbf{u}$ is the velocity vector, T is the temperature, p is the non-lithostatic pressure, $\hat{\mathbf{k}}$ is the unit vector in the direction of gravity and t is the time. In addition, the dimensionless parameter:

$$Ra = \frac{\beta g \Delta T d^3}{\kappa \nu} \quad (3.4)$$

denotes the Rayleigh number, where g is the acceleration due to gravity, β is the coefficient of thermal expansion, ΔT is the temperature drop across the domain, d is the domain depth, κ is the thermal diffusivity and ν is the kinematic viscosity.

A widely used two-dimensional (2-D) geodynamics finite element program, CONMAN, which employs quadrilateral elements and bilinear shape functions for velocity, is utilized to solve these incompressible, infinite Prandtl number equations. The main characteristics of the code are presented here, although a more detailed description can be found in King et al. (1990) and Appendix B. The continuity equation is treated as a constraint on the momentum equation and incompressibility is enforced in the solution of the momentum equation using a penalty formulation. The well known streamline upwind Petrov Galerkin (SUPG) method is used to solve the energy equation (Hughes & Brooks 1979) and an explicit second order predictor corrector algorithm is employed for time marching. Since the temperatures provide the buoyancy to drive the momentum equation and, as there is no time dependence in the momentum equation, the algorithm to solve the system is simple: given an initial temperature field, calculate the resulting velocity field. Use the velocities to advect the temperatures for the next time step and solve for a new temperature field.

3.3.2 Adaptive Strategies

Over recent decades, unstructured grid systems have been developed and applied in simulations of various computational fluid mechanics problems. The accuracy of a computational

solution is strongly influenced by the discretization of the space in which a solution is sought. In general, the introduction of a highly dense distribution of nodes throughout the computational domain will yield a more accurate answer than a coarse distribution. However, limitations in computer processing speed and accessible memory prohibit such a scenario. An appropriate alternative would be to improve the accuracy of the computation where needed. Grid adaptation provides a suitable means to do this, ensuring that grids are optimized for the problem under study. Broadly speaking, such adaptive procedures fall into two categories:

1. h -refinement, in which the same class of elements continue to be used, but are changed in size to provide the maximum economy in reaching the desired solution.
2. p -refinement, in which the same element size is utilized, but the order of the polynomial is increased or decreased as required.

A variant of the h -refinement method, termed adaptive remeshing, is employed in this study. It provides the greatest control of mesh size and grading to better resolve the flow features. The main advantages offered by such methods are (Lohner 1995):

1. the possibility of stretching elements when adapting features that are of lower dimensionality than the problem at hand, which leads to considerable savings;
2. the ability to accommodate, in a straightforward manner, problems with moving bodies or free surfaces.

In this method, the problem is solved initially on a coarse grid, noting that this grid must be sufficiently fine to capture the important physics of the flow. Remeshing then involves the following steps:

1. The solution on the present grid is analyzed through an error indication procedure, to determine locations where the mesh fails to provide an adequate definition of the problem. An interpolation based local error indicator is employed in this study, based upon nodal temperature curvatures (e.g. Peraire et al. 1987).

2. Given the error indication information, determine the nodal spacing, δ , the value of the stretching parameter, s , and the direction of stretching, α , for the new grid.
3. Using the old grid as a background grid, remesh the computational domain utilizing a variant of the advancing front technique (George 1971, Lo 1985, Peraire et al. 1987, Davies et al. 2007), which is capable of generating meshes that conform to a user prescribed spatial distribution of element size (i.e. δ , α and s).
4. Interpolate the original solution between meshes.
5. Continue the solution procedure on the new mesh.

The remeshing process is repeated until the desired solution criteria are met (see Appendices C and D for further details).

3.4 Geodynamical Applications

3.4.1 Mid–Ocean Ridge Magmatism

A significant body of work has been published on the numerical modeling of Mid–Ocean Ridges (MOR). For example, Buck et al. (2005) use numerical models to study modes of faulting at ridges, Kuhn & Dahm (2004) employ numerical models to study magma (i.e. molten rock) ascent beneath ridges, while Albers & Christensen (2001) study the channeling of plumes below ridges. While these models were designed to investigate specific processes at ridges, incorporating complex material properties and boundary conditions, we present a simple, generic, passive (buoyancy forces are neglected) MOR model, utilizing our results to demonstrate the benefits of grid adaptivity.

Model Geometry and Boundary Conditions

The model presented does not incorporate the entire convecting mantle. Instead, we focus on the region directly adjacent to a MOR. Our results are derived from simulations in a




Figure 3.1: A summary of the boundary conditions utilized in our mid-ocean ridge model, where MBC denotes the mechanical boundary conditions. Stress free conditions, i.e. no normal or shear stress, are employed at the lower and side boundaries, with prescribed velocities (kinematic) on the upper boundary, the non-dimensional equivalent of 5cm yr^{-1} . Temperatures are fixed on upper ($T = 0$) and lower ($T = 1$) boundaries with insulating sidewalls.

rectangular domain of height 1, which in our application to a MOR represents 500km, and width 5 ($= 2500\text{km}$), x being non-dimensionalized (x') as:

$$x' = \frac{x}{l_0} \quad (3.5)$$

where $l_0 = 500\text{km}$. By limiting the vertical and horizontal extent of the domain, computational expenditure is reduced, allowing one to accurately resolve the flow details contiguous to plate boundaries. The main drawback of this technique is that flow must be permitted through the lower and side boundaries of the model; these boundary conditions are therefore specified in such a way as to mimic the effect of the full convecting system on the smaller region under study (Figure 3.1).

Plate motion is prescribed as a kinematic boundary condition at the upper surface. A non-dimensional velocity, equivalent to 5cm yr^{-1} , is chosen, utilizing the non-dimensional




Figure 3.2: The thermal field generated by our mid-ocean ridge simulations. Red is hot ($T = 1$), blue is cold ($T = 0$) and the color scale is linear. A series of stream-traces are included, indicating the flow field behavior.

relation:

$$v' = \frac{vl_0}{\kappa} \quad (3.6)$$

where κ denotes thermal diffusivity, taken as $1 \times 10^{-6} \text{ m}^2\text{s}^{-1}$. Time is non-dimensionalized by the conductive time scale:

$$t' = \frac{t\kappa}{l_0^2} \quad (3.7)$$

The model makes no attempt to account for forces that move the plate. The new plate that is continuously created within the model is disposed of by a prescribed rate of flow through the outer boundary. This in turn is replaced by material from the side and lower boundaries. Material properties are uniform throughout the domain, there are no internal heat sources and the Rayleigh number is set to zero.

Results

We find that the broad patterns observed in previous studies are reproduced (Figure 3.2). With flow driven kinematically by surface ‘plate’ motion, the system evolves to approach

a steady state. This typically involves the development of a cold ‘plate’ thickening with age at the surface, with flow beneath focusing heat directly towards the ridge axis (i.e. the upper center of our model). This ‘plate’ is particularly well captured in our simulations, as a direct consequence of the adaptive methodologies utilized. It is necessary, therefore, to provide a brief run through the evolution of the calculation, to illustrate the benefits of grid adaptivity.

Having obtained an initial solution on a coarse grid (Figure 3.3–Stage 1b), mesh adaptation was invoked to resolve, in more detail, the temperature profile encountered. The solution was analyzed via the error indication procedure and the domain remeshed, utilizing the information yielded by this error indicator (the generation parameters δ_{Min} , δ_{Max} , s_{Max} and C displayed in Table 3.1) to control the regeneration process. The ensuing grid is displayed in Figure 3.3 (Stage 2a). Note that nodes have automatically clustered around zones of high temperature gradient, at the surface and immediately below ‘plate’ boundaries.

The solution procedure continued on this new mesh, producing the thermal field observed in Figure 3.3 (Stage 2b). It is clear, even visually, that the solution on this grid is far better resolved than that seen in Figure 3.3–Stage 1b, with contours more steady and consistent. However, by examining the thermal field close to the ridge axis (Figure 3.4–Stage 2b) it becomes apparent that the problem remains inadequately defined. Consequently, one further remeshing loop was invoked, producing the mesh observed in Figure 3.3 (Stage 3a). The simulation was terminated once the solution was deemed to have converged on this mesh. The final temperature profile is displayed in Figures 3.3 and 3.4 (Stage 3b).




Figure 3.3: Evolution of the temperature field (b) on a series of adapted grids (a). Red is hot ($T = 1$), blue is cold ($T = 0$), and the contour spacing is 0.05. Note that the coordinate scales are distorted, with $x : y = 0.5$.




Figure 3.4: As in Figure 3.3, but at higher magnification, close to the ridge axis. Once again, red is hot ($T = 1$), blue is cold ($T = 0$), and the contour spacing is 0.05

Stage	Elements	Nodes	δ_{min}	δ_{max}	s_{max}	C
1	4096	4225	-	-	5	-
2	12965	13259	0.01	0.1	5	0.05
3	22790	23263	0.005	0.1	5	0.03

Table 3.1: The sequence of meshes and mesh generation parameters employed for the mid-ocean ridge simulations. It should be noted that the initial mesh (Stage 1) was generated via a simple uniform mesh generator, as opposed to the advancing front generator typically employed throughout this study.

With each remeshing, the benefits of the multi-resolution solution permitted by the adaptive methodologies can be appreciated. Within the upper thermal boundary layer, where temperature contours are extremely compact and gradients are high, a large number of nodes is required to generate an accurate solution. Since the lower reaches of the model are more passive, with reduced solution gradients, the number of nodes required for accuracy is significantly less. The proposed method automatically ensures that an ‘optimal’ mesh is generated, with zones of fine resolution being analogous to zones of high solution gradient. Consequently, the thermal field can be adequately resolved.

As a quantitative test of the method, we have computed heat flow as a function of ocean floor age, at each stage of the calculation (i.e. for converged solutions on each grid). The results are displayed in Figure 3.5, alongside data derived from a cooling half-space model (Turcotte & Schubert 2002) which is an analytical approximation to the problem, and data obtained from a simulation utilizing a fully uniform, structured mesh of almost 30,000 elements (SM in Figure 3.5). An exceptional agreement is observed between all data sets beyond $\approx 1\text{Myr}$ (Myr = million years). This accord, however, disappears within $\approx 1\text{Myr}$ of the ridge axis. Here, the half-space model tends towards infinity, whereas our simulations converge towards a finite value, as indeed would be expected from the physics of the numerical problem. Nevertheless, this graph provides a simple way to illustrate the benefits of the proposed methodology. Simulation results show sequentially improving agreements with the half-space model as one refines the grid from Stage 1 to Stage 3. Results track the




Figure 3.5: Plot of computed heat flow against age relation for our MOR models and the cooling half-space model for $k = 3.3 \text{ Wm}^{-1}\text{K}^{-1}$. A clear divergence is observed between the computed solution and the cooling half-space. However, as the grid becomes more and more refined, through the adaptive procedures, this divergence decreases dramatically. Note that SM represents the solution obtained on a fully uniform structured mesh. It is included for comparative purposes only.

analytical solution closer to the ridge axis, culminating in successively higher values at the axis itself. This is a direct consequence of the improved resolution inherent to the adaptive methodologies utilized. At the ridge axis itself, the true numerical solution is extremely difficult to reproduce. However, what is clear from this graph is that the adaptive methodologies employed significantly improve solution quality. The results from a fully uniform structured mesh, albeit with more degrees of freedom, are not competitive with those obtained using the adapted grids.

In addition to increasing solution accuracy, the adaptive refinement strategies are computationally highly efficient. As regards to the MOR simulations presented here, for a specified level of accuracy a uniform mesh simulation expends approximately 20% more CPU time than an adaptive mesh simulation, with figures for the adaptive case including the time allocated for remeshing. The generation of a new optimal mesh is an inexpensive procedure,

typically taking between 15 and 20 time-steps, compared to the time taken for one time step with a fixed mesh. It should be noted however that the time expended in remeshing can be decreased significantly by specifying a larger remeshing tolerance, *msh_tol*.

3.4.2 Subduction Zone Magmatism

As is the case with mid-ocean ridges, numerical models have become central in shaping our understanding of SZ dynamics and thermal structures. Andrews & Sleep (1974) use numerical models to demonstrate that frictional heating along the subducting plate is not likely to produce enough heat to melt the slab. Davies & Stevenson (1992) cite numerical models as primary evidence to suggest that the oceanic crust of the down going slab is not melted extensively, if at all, and, hence, is not the source of SZ magmatism, with the possible exception of the special case of very young oceanic crust, which is hotter. Numerical simulations have also been developed for studies of SZ mineralogy and metamorphism (Peacock 1996), transportation of water and its influence on melting (Iwamori 1998), the thermal and dynamic evolution of the upper mantle in SZ (Kincaid & Sacks 1997) and the effects of chemical phase changes on the downwelling slab (Christensen 2001). It is important to note that the SZs discussed here have an idealized geometry. The terminology used in describing such SZs is presented in Figure 3.6.

In this study, a basic isoviscous flow model is presented which is used to demonstrate the benefits of grid adaptivity within a SZ context. It is common knowledge that the most difficult region to resolve in any SZ model is the area between the subducting slab and the overriding plate, commonly known as the mantle wedge corner, as a direct consequence of a singularity at the intersection between slab and plate. Since the most geologically important processes in SZ occur here, models of wedge flow need to be carefully constructed. Previous studies have achieved higher resolution in this area by *a priori* generating a mesh with a




Figure 3.6: The geometry of a generic subduction zone. Note that by overriding plate we mean the rigid lithosphere. By wedge corner we mean the apex at which the overriding plate and slab meet. Adapted from Davies & Stevenson (1992).

large number of nodes clustered in the wedge (e.g. Davies & Stevenson 1992). However, this is not ideal. Grid adaptivity provides a suitable alternative, allowing one to automatically generate an optimal mesh, through *a posteriori* error indication procedures, ensuring that nodes are positioned where required. Such techniques could therefore play an important role in future solution strategies for these models, for both the steady state simulations considered here and more complex unsteady problems.

Model Geometry and Boundary Conditions

We do not simulate the entire convecting mantle. Instead, we focus on the region directly adjacent to a generic SZ. The results of calculations using a box of 3×2 non-dimensional units are presented, which is equivalent to 600×400 km. Boundary conditions are summarized in Figure 3.7. We shall distinguish two lithospheres. First, a mechanical lithosphere, which will be considered to be the rigid part of the plate on the time scale of the process, and, second, a thermal lithosphere, which is Earth's upper thermal boundary layer. The descending slab is prescribed by a kinematic boundary condition as a non-migrating slab dipping uniformly at 60° . The subduction velocity is set to the non-dimensional equivalent




Figure 3.7: A composite diagram illustrating the boundary conditions utilized in the subduction zone model.

of 9 cm yr^{-1} . These velocities are also prescribed to the incoming plate. A zero velocity condition is specified at certain nodes to model the mechanical lithosphere of the overriding plate, corresponding to a thickness of 50km. This restricts them from participating in the viscous flow region. The thickness of the mechanical lithosphere of the downgoing plate is taken to be the same as the overriding plate (i.e. 50 km). The side and lower domain boundaries are prescribed with velocities derived from the analytical solution to a Newtonian corner flow problem (McKenzie 1969). Indeed, by setting the model up in this way, a direct comparison can be made between simulated velocities and those of the analytical solution. This allows a quantitative demonstration into the benefits of grid adaptivity to SZ simulations.


Temperature boundary conditions are slightly more complex. The temperature is fixed at the surface ($T = 0$) and a zero heat flux condition is specified at the base of the model. On the continental side, i.e. the overriding plate, the thermal boundary layer is assumed

to be 100km thick. Within this layer, it is assumed that vertical heat transfer is practically by conduction alone and that steady state conditions prevail. The temperature profile is, therefore, represented by a linear temperature gradient, with the temperature at the base of this layer assumed to be 1350°C, or $T = 1$ in non-dimensional units, and the temperature at the top, i.e. the surface of the overriding plate, assumed to be 0°C, or $T = 0$ in non-dimensional units. The situation on the oceanic side, i.e. the incoming plate, is slightly different. The oceanic plate is created at the axis of a mid-ocean ridge and cools as it moves away from the ridge axis, as was illustrated in Figure 3.2. The temperature profile within the incoming plate can, therefore, be approximated, by a standard error function, consistent with a plate of age 40Ma (Carslaw & Jaeger 1959), as:


$$T(y) = T_{(s)} + \left\{ T_{(m)} - T_{(s)} \right\} \operatorname{erf} \left\{ \frac{y}{2\sqrt{\kappa t_{40}}} \right\} \quad (3.8)$$

where t_{40} is the age of the plate in seconds, $T_{(s)}$ is the surface temperature, $T_{(m)}$ is the mantle temperature and κ is the thermal diffusivity, which is assumed to be $10^{-6} \text{m}^2 \text{s}^{-1}$. The thermal boundary layer is assumed to be 100 km thick with temperatures at its top and base taken as $T = 0$ and $T = 1$, in non-dimensionalized units, respectively.

Convection is believed to be the dominant mode of heat transfer in the upper mantle, beneath the thermal boundary layers. Consequently, the temperature gradient is lower. In the interior of a vigorously convecting fluid, the mean temperature gradient is approximately adiabatic. Considering the adiabatic temperature gradient of the uppermost mantle (Turcotte & Schubert 1982) and the depth of the box, i.e. 400km, a temperature of 1470°C is specified at the bottom left and right corners of the box, which, in non-dimensionalized units, is ≈ 1.09 . The temperature increases linearly from the base of the thermal boundary layer to this point. There are no explicit heat sources or sinks within the model.



(a)



(b)

Figure 3.8: (a) The steady state thermal field yielded by our subduction zone simulations. Red is hot ($T = 1.09$), blue is cold ($T = 0$) and the color scale is linear. Note that the slab remains cool throughout, while the mantle wedge corner heats up significantly. The dark lines traversing the solution domain are velocity stream-traces, included to provide an indication of the direction of motion. (b) The final mesh.

Results

Results are displayed in Figure 3.8. They are broadly consistent with previous SZ models, with the thermal field being characterized by rapid temperature variations over the solution domain, predominantly along the margins of the subducting plate and in Earth's upper thermal boundary layer. However, these results are not central to our study. It can be seen from Figure 3.8 that the adaptive procedure has refined the grid at locations of high temperature gradient, without overloading the remainder of the domain (note that a preset element size is specified for nodes in the upper mechanical lithosphere and the down going slab, since velocities here are prescribed). Such grid refinement has a dramatic effect on solution accuracy. This is illustrated in Figures 3.9 and 3.10, which display the discrepancy between simulated velocities and those yielded by the analytical solution. This local error, E_L , is calculated as:

$$E_L = \frac{|V_M - V_A|}{|V_A|} \quad (3.9)$$

where V_M denotes simulated velocities, V_A the velocities yielded by the analytical solution and $|\cdot|$ absolute value. The improvements yielded by grid adaptivity are clear to see. On the initial mesh (Stage 1), the error is extremely prominent, emanating from its source at the mantle wedge corner and strongly degenerating the solution over a large section of the wedge. A minor error is also visible at the corner underlying where the incoming plate descends to become the down going slab, although it is small in comparison to that observed in the wedge corner. By Stage 2, the re-meshing process has refined the grid considerably in these regions and, consequently, a substantial decrease in error is observed. An additional reduction in error also occurs in Stage 3, as the grid becomes further refined at these locations. Even though the effects of the singularity are not fully nullified, its

Stage	Elements	Nodes	δ_{min}	δ_{max}	s_{max}	C
1	7333	7499	0.06	0.06	1	-
2	17876	18131	0.02	0.04	3	0.2
3	20906	21157	0.01	0.04	3	0.1

Table 3.2: The sequence of meshes and mesh generation parameters employed for the subduction zone simulations

influence is severely restricted by the grid refinement procedure.

The point is reinforced by examining the mean global error, E_G , calculated as:

$$E_G = \int_{\Omega} E_L d\Omega / \int_{\Omega} d\Omega \quad (3.10)$$

Results are presented in Table 3.3, demonstrating quantitatively that the refinement process undoubtedly has a positive influence on the global error. As the grids are adapted a dramatic decrease in error is observed. This is particularly true for the first remeshing, where E_G decreases by a factor of 4.5.

In summary, the adaptive strategies employed have significantly improved solution accuracy for the SZ simulations presented here. The refinement process has severely restricted the influence of the intersection singularities and, consequently, solution accuracy throughout the domain is improved. Results suggest that an extension of this work to models with more realistic mantle rheologies, i.e. material properties, together with more Earth like surface plate behavior, incorporating solidification and localization phenomena, would be a worthwhile exercise. A true understanding of this system will only be gained by studying coupled crustal/mantle models. Such models naturally require fine resolution within the crust, where the plates fracture, bend and buckle, and coarser resolution as one descends into the mantle, where deformation occurs on a much larger scale. Error-guided grid adaptivity should therefore be an invaluable tool in simulating such dynamic systems.




Figure 3.9: (a) The series of adapted grids, (b) The discrepancy between simulated and analytical velocities, E_L , (c) A high resolution image of this error in the mantle wedge corner. Contour values are not important in this figure; they are included in Figure 3.10. This figure merely illustrates the extent of the error decreases with each remeshing. The influence of the singularity at the mantle wedge corner is severely restricted by the grid refinement procedure.




Figure 3.10: High resolution contour plots of the solution error, E_L , in the mantle wedge corner. Contour values range between 0.1 and 0.9, at a contour spacing on 0.2. (a) represents the solution error obtained on the initial grid, (b) the error after 1 remeshing, while (c) represents the final error, i.e. after 2 remeshings.

Stage	E_G
1	0.0117
2	0.0026
3	0.0022

Table 3.3: The mean global error, E_G , obtained following each stage of the subduction zone simulations.

3.5 Conclusions

An adaptive finite element procedure has been applied in simulations of two separate geodynamical processes—fluid flow at a mid-ocean ridge (MOR) and at a subduction zone (SZ). The method has refined the locations of thermal boundary layers wherever they are strong, at the ridge itself and along Earth’s surface (MOR), and in the mantle wedge, along the margins of the descending plate and at Earth’s surface (SZ). The adapted grids maintain good solution accuracy and, through a series of remeshings, display the ability to gradually improve solution quality, without significantly increasing the total number of unknowns at each stage. The advocated methods are computationally highly efficient, expending approximately 20% less CPU time than uniform mesh simulations, for a specified level of accuracy.

This investigation suggests that coupling adaptive strategies to more complex models will lead to a new class of geodynamical simulation, yielding greater insights into the intricate processes at work within Earth’s interior. With the methods currently employed in the field, such insights are beyond our capabilities. However, memory efficient numerical techniques, such as the adaptive strategies presented here, will ensure that research is not unnecessarily restricted by computer power. It is therefore of great importance that the geodynamical community begins to implement such schemes.

Chapter 4

Towards Automatic

Multi-resolution in 3-D Spherical

Models of Mantle Convection

4.1 Abstract

A method for introducing multi-resolution capabilities into pre-existing 3-D spherical mantle convection codes is presented. The method, based upon the multigrid refinement techniques employed in the field of computational engineering, is utilized to refine and solve on a radially non-uniform grid. However, it is equally applicable for complete spatial refinement, i.e. radially and/or laterally on a sphere. Highly efficient solutions to non-uniform problems are obtained, consisting of a high resolution upper mantle with coarser resolution elsewhere. Since the viscosity of Earth's mantle increases significantly with depth, upper mantle convection is dominated by fine-scale activity, with longer wavelength features more prevalent in the lower mantle. The advocated scheme fits in perfectly with such behavior, yielding

great benefits for geodynamical investigations. It is highly efficient in terms of RAM, meaning that one can attempt calculations that would otherwise be impractical. In addition, the scheme reduces the CPU-time needed to solve a given problem. Validation tests illustrate that the method is accurate and robust.

4.2 Introduction

4.2.1 Motivation and Background

Plate tectonics is a theory that underpins most of modern geology, with the basic idea being that Earth's surface consists of several rigid plates that are in relative motion with respect to one another. While sea floor spreading and subduction are widely accepted mechanisms for the creation and destruction of these plates, the underlying processes, which are responsible for the observed pattern of plate motion, are not comprehensively understood. The general consensus is that plate movement is the surface expression of a global pattern of solid state convection within Earth's mantle. Indeed, in some way or another, ultimately, all large-scale geological activity and dynamics of the planet can be related to the release of potential energy within the mantle.

Earth's present-day mantle is obscured from direct view. Nevertheless a wealth of information has accumulated regarding its structure, composition and internal dynamics. Seismology (e.g. Dziewonski et al. 1975, Kennett & Engdahl 1991, Grand et al. 1997, van der Hilst et al. 1997, Bijwaard & Spakman 1999, Kennett & Gorbato 2004), accurate measurements of the gravity and geoid field (e.g. Richards & Hager 1984, Ricard et al. 1993, Tapley et al. 2005) and heat flow studies (e.g. Bullard 1939, Jessop et al. 1976, Pollack et al. 1993, Stein & Stein 1994) have all played a role in this development. This observational knowledge has been supplemented by geochemical and mineral physics experiments, some conducted

at extreme conditions of temperature and pressure (e.g. Allegre et al. 1996, Hoffman 1997, Bercovici & Karato 2003, Murakami et al. 2004, Oganov & Ono 2004, Boyet & Carlson n.d.). A multidisciplinary team of scientists has therefore come together, with a mutual aspiration to understand Earth's mantle.

An imperative component in this multi-faceted approach is the ability to numerically simulate the mantle's convective flow field. This is best achieved within global, three-dimensional (3-D) spherical models which solve the conservation equations of mass, momentum and energy, at infinite Prandtl number, in a spherical shell (e.g. Baumgardner 1985, Bunge et al. 1996, Moresi & Solomatov 1995, Zhong et al. 2000, Stemmer et al. 2006, Choblet et al. 2007). The existing 3-D spherical codes, e.g. TERRA (Baumgardner 1985), CITCOMs (Zhong et al. 2000) and OEDIPUS (Choblet et al. 2007), are computationally highly efficient, utilizing an almost uniform discretization of the sphere and resourceful solution strategies, frequently centered around a multigrid solver. They have shown their value on countless occasions, yielding otherwise unattainable information on the temporal and spatial evolution of mantle flow under a range of rheologies and boundary conditions (e.g. McNamara & Zhong 2005, Davies 2005, Davies & Bunge 2006, Zhong 2006) and also, by predicting the seismic tomography anomalies found in the mid-mantle (Bunge & Grand 2000, Bunge & Davies 2001). However, since mantle convection is nonlinear and time-dependent, these codes are difficult to formulate and validate. Furthermore, the scale and geometry of the problems under investigation mean that, regardless of their numerical efficiency, even the simplest simulations are computationally highly demanding. As a consequence, despite the massively parallel computational architectures available to us today, it is often impossible to resolve thermal and chemical boundary layers at 'Earth-like' Rayleigh numbers. This task becomes even more difficult for 'early Earth' simulations, likely to be characterized by more vigorous convection. This is a major concern.

4.2.2 The Need for Multi-Resolution

Mantle convection is characterized by fine scale upwelling and downwelling activity in a more passive, large-scale background flow. Such narrow, high gradient regions present a serious challenge for computational methods: their location and extent is often undeterminable and, even if they are identified, with current methods it is often impossible to resolve localized features. Indeed, one of the most challenging problems currently facing geodynamicists is the accurate solution of such multi-scale flows.

In finite element analysis, non-uniform resolution is usually attained via unstructured grids, with solution accuracy and computational efficiency improved through error-guided grid adaptivity (e.g. Peraire et al. 1987, Nithiarasu & Zienkiewicz 2000, Davies et al. 2007). As illustrated in Chapters 2 and 3, there are clearly significant advantages to such techniques. Indeed, when compared to uniform grids, for the same, or often superior precision, the number of degrees of freedom is reduced when the adaptive, unstructured procedures are employed. This leads to an increase in computational efficiency, in terms of both processing time and memory demands. Additionally, the flexibility warranted by unstructured grids makes discretizing many of the complex geometries encountered on Earth a simple task, particularly if triangles are used in two-dimensions (2-D) and tetrahedra in 3-D.

Occasionally however, such techniques lead to more complication than is necessary. In view of the existing 3-D spherical mantle dynamics codes, which, as stated previously, are built upon a quasi-uniform discretization of the sphere and closely structured solution strategies, such techniques would throw away the regular grid and, with it, the major benefits of the current solution algorithms. Alternative avenues towards non-uniform resolution must be sought. A method capable of producing similar advantages to unstructured grids is needed, whilst maintaining the key benefits of the current configuration and also being

straightforward to implement. A multigrid refinement, or ‘adaptive multigrid’, approach provides the answer.

The multigrid method (e.g. Brandt 1984, Briggs et al. 2000, see Appendix F for further details) is an amalgamation of ideas and techniques, combining iterative solution strategies and a hierarchy of discretizations/grids, to form a very powerful tool for the numerical solution of differential equations. The basic idea behind the technique is to work not with a single grid, but with a sequence of grids (‘levels’) of increasing fineness. The method can be applied in combination with any of the common numerical discretization techniques and, consequently, has been widely used within the geodynamical community (e.g. Baumgardner 1985, Tackley 1996, Bunge et al. 1997, Zhong et al. 2000, Kameyama et al. 2005, Choblet et al. 2007). The ‘adaptive multigrid’ can be considered a multigrid process in which the fine grid covers only part of the domain. It is the limited extent of the fine grid that provides the benefits to the method.

The addition of finer mesh components spanning only parts of the domain was described early in multigrid history (Brandt 1977), but the idea has only occasionally been put into practice (e.g. Bai & Brandt 1987, Thompson et al. 1992, Rude 1993, Lopez & Casciaro 1997, Ekevid et al. 2004).

The applicability of such techniques to 3-D spherical infinite Prandtl number codes is investigated here. Our main aim is to demonstrate that these solution algorithms can be successfully incorporated within pre-existing codes, in a straightforward manner. The chapter begins with a general overview of the adaptive multigrid technique, utilizing a two-dimensional 2-D rectangular domain to illustrate the key ideas. The implementation of a comparable scheme within a pre-existing 3-D spherical mantle dynamics code is then outlined, in the context of TERRA (Baumgardner 1985, Bunge & Baumgardner 1995). This

implementation is termed ‘radial refinement’; refinement is restricted to the radial component of the sphere, being achieved via the adaptive multigrid methodology. The scheme, and consequently TERRA, is then validated by (i) solving thermal convection problems with known benchmark solutions and (ii) comparing results with quasi-analytical solutions derived via propagator matrix methods. Results indicate that the proposed methodology is highly successful, maintaining solution accuracy whilst improving computational efficiency.

4.3 Methodology

4.3.1 Multigrid

Multigrid is an iterative method for the solution of algebraic systems arising from discrete approximations to ordinary or, more commonly, partial differential equations. The idea originated from attempts to improve the slow convergence of classical iterative (formally termed relaxation) methods. Consider a linear algebraic equation of the form:

$$Au = f \tag{4.1}$$

where the solution u is sought, with the size of the matrix A being extremely large. Since direct matrix inversion is out of the question for such large systems, an iterative approach is commonly employed. Standard iterative methods, such as Jacobi and Gauss-Seidel, are known to be very effective for the first few iterations. Inevitably however, convergence slows and the schemes appear to stall. Furthermore, the rate of convergence reduces as the problem size increases. This is because such schemes possess the so-called *smoothing property*: they are very effective at eliminating high-frequency or oscillatory components of the error, while leaving low-frequency or smooth components relatively unchanged (e.g. Brandt 1984, Briggs et al. 2000).

Relaxation, therefore, efficiently reduces non-smooth error components, leaving a relatively smooth error, which can be well-approximated on a coarser grid, where solution is much cheaper. A caveat to this is that fine-grid error modes appear more oscillatory on a coarse grid. This suggests that when relaxation begins to stall, signalling the predominance of smooth error modes, it is advisable to move to a coarser grid; there, the smooth error modes appear more oscillatory and relaxation will be more effective. Multigrid methods exploit the smoothing property of iterative methods to improve the rate of convergence by utilizing coarser grids. The question now is: how does one move to a coarser grid and relax on the more oscillatory error modes?

It is at this point that the multigrid technique begins to come together. Relaxation sweeps very quickly reduce all high-frequency components of the error. Its smoother part should then be reduced by being approximated on a coarser grid, a grid with mesh-size $H = 2h$, for example, where h is the mesh-size of the finer grid. Generally, for any linear fine-grid equation $A^h u^h = f^h$ and any approximate solution v^h , the error, $e^h = u^h - v^h$, satisfies the *residual equation*:

$$A^h e^h = r^h \quad \text{where} \quad r^h = f^h - A^h v^h \quad (4.2)$$

This error, e^h can therefore be approximated by the coarse-grid function e^H , which satisfies:

$$A^H e^H = I_h^H r^h \quad (4.3)$$

where A^H is some coarse-grid approximation to A^h , and I_h^H is a fine-to-coarse transfer operator, more commonly known as a *restriction* operator. In other words, $I_h^H r^h$ is a coarse-grid function whose value at each point is a certain weighted average of values of r^h at neighboring fine-grid points.

Having obtained an approximate solution e^H to Equation 4.3, we use it as a correction to

the fine-grid solution. Namely, we replace:

$$v_{new}^h \leftarrow v^h + I_H^h e^H \quad (4.4)$$

where I_H^h is a coarse-to-fine *interpolation* or *prolongation*. This means that at each fine-grid point, the value of $I_H^h e^H$, designed to approximate the error e^h , is interpolated from values of e^H at neighboring coarse-grid points. Linear interpolation can be used in most cases. An important point to note here is that interpolation is most effective when the error is smooth. If the error is oscillatory, even a very good coarse-grid approximation may produce an interpolant that is not very accurate. Fortunately, this provides a complement to relaxation, which is most effective when the error is oscillatory.

The whole process of calculating $I_h^H r^h$, solving Equation (4.3) and interpolating the correction (Equation 4.4) is called a *coarse grid correction*. This procedure is the basis of the so-called *correction scheme* (CS). Having relaxed on the fine grid until convergence deteriorates, we relax the residual equation on a coarser grid to obtain an approximation to the error itself. We then return to the fine grid to correct the original approximation. It is important to appreciate how well the various elements complement each other during the whole process. Relaxation on the fine grid eliminates the oscillatory components of the error, leaving a relatively smooth error, which can be accurately transferred to a coarser grid. Assuming the residual equation can be solved accurately on a coarse grid, it is still important to transfer the error accurately back to the fine grid. Because the error is smooth, interpolation will work very well and the correction to the fine-grid solution should be effective.

The aim of multigrid methods is to attain a rate of convergence which is independent of the number of grid points, N , so that the total amount of work involved in obtaining a solution is proportional to the number of grid-points. The two grid scheme, as outline above, should

converge faster than the basic one-grid method, however, it will not achieve the goal of a multigrid scheme. If the system of coarse grid equations is solved by some direct method, rather than an iterative method, the rate of convergence becomes independent of N , but the work involved in the direct solution grows rapidly with N , and will normally become prohibitive for large N . If a relaxation method is used, the rate of convergence on the coarse grid is better than on the fine grid, but it still depends on N .

Multigrid methods achieve the required result through solving the coarse-grid equations by invoking the same process recursively; a few smoothing iterations are performed on the coarse grid, the new residuals are restricted to an even coarser grid, the equations are solved on this grid etc... The method, therefore, exploits a succession of grids, each one being coarser than the previous. Continuing the nesting process will eventually lead to a grid with such a small number of points that a direct solution is feasible, or such that an iterative process will converge rapidly.

4.3.2 The Full Approximation Storage Scheme

The *Full Approximation Scheme* (or Full Approximation Storage - FAS) is a widely used version of multigrid processing. It has mainly been used in solving non-linear problems, but it has so many other applications that it should perhaps be used in most advanced programs. In the FAS Scheme, exactly the same steps are performed as in the CS, but in terms of another coarse-grid variable. Instead of e^H , we use:

$$u^H = I_h^H v^h + e^H \quad (4.5)$$

as the coarse-grid unknown function. This coarse-grid variable u^H approximates $I_h^H u^h$, the full intended solution represented on the coarse grid and, hence, the name 'Full Approximation Scheme'. The FAS coarse-grid equations, derived from Equations (4.3) and (4.5)

are:

$$A^H u^H = \hat{f}^H \quad (4.6)$$

where:

$$\hat{f}^H = A^H(I_h^H v^h) + I_h^H r^h \quad (4.7)$$

Having obtained an approximate solution, $\tilde{u}^H$, to 4.6, the approximate coarse-grid correction is, of course:

$$e^H = \tilde{u}^H - I_h^H v^h \quad (4.8)$$

The FAS interpolation back to the fine grid, equivalent to Equation (4.4) is:

$$v_{new}^h \leftarrow v^h + I_H^h(\tilde{u}^H - I_h^H v^h) \quad (4.9)$$

Note that, to use directly:

$$v_{new}^h \leftarrow I_H^h \tilde{u}^H \quad (4.10)$$

would be worse as it introduces the interpolation errors of the full solution u^H instead of the interpolation errors of only the correction, e^H .

It is important to note that for linear problems, the FAS steps are fully equivalent to the CS steps. However, in the FAS mode, rather than store the correction at each grid level (designed to correct the finer-level approximation), the idea is to store the full current approximation, which is the sum of the correction and its base approximation. This is an essential prerequisite to any refinement strategy based around a multigrid, since in parts of the domain not covered by the finest grid, the coarser grid must show the full solution, not just a correction.

Equivalence of CS and FAS Schemes for Linear Problems

For linear problems, the FAS steps are fully equivalent to the CS steps. Taking Equation (4.6) and substituting for Equation (4.7) gives:

$$A^H \tilde{u}^H = A^H (I_h^H v^h) + I_h^H r^h \quad (4.11)$$

Multiplying both sides by $(A^H)^{-1}$ yields:

$$\tilde{u}^H = (A^H)^{-1} [A^H (I_h^H v^h) + I_h^H r^h] \quad (4.12)$$

which, when A is a linear operator, reduces to:

$$\tilde{u}^H = I_h^H v^h + (A^H)^{-1} (I_h^H r^h) \quad (4.13)$$

Rearranging this equation gives:

$$\tilde{u}^H - I_h^H v^h = (A^H)^{-1} (I_h^H r^h) \quad (4.14)$$

and utilizing Equation (4.8):

$$e^H = (A^H)^{-1} (I_h^H r^h) \quad (4.15)$$

which is identical to the coarse-grid corrections obtained via the CS (Equation 4.3). In other words, for linear problems, the FAS interpolation back to the fine grid (Equation 4.9) simplifies to the CS fine-grid corrections displayed in Equation (4.4). Both schemes are directly equivalent.

4.3.3 The Adaptive Multigrid

Multigrid programs commonly use uniform grids at each level of refinement, as illustrated by Figure 4.1. Such a configuration makes programming more straightforward, whilst also avoiding the computational overhead of dealing with varying mesh spacing. Of course, in




Figure 4.1: An example of the hierarchy of uniform grids used in regular multigrid cycles. A standard bisection refinement rule is employed; each quadrilateral element is split into four elements at the next grid level. Each grid is labeled with the number of nodes (in italic) and the number of elements (in bold).

many situations, uniform grids lead to an excessive problem size and, consequently, models that are computationally inefficient and impractical. Fortunately, multigrid algorithms need not necessarily be restricted to truly uniform discretizations. The approach described here is to recover the flexibility of non-uniform grids by observing that the various grids used in the usual multigrid cycles need not extend over the same domain. The finest levels may be confined to progressively smaller subdomains, so as to provide higher resolution exactly where desired. These ‘local patches’ are treated identically to ‘global’ grids in the multigrid algorithm, only that their boundary values are obtained by interpolation from coarser grids, where needed.

The majority of numerical solutions are derived in two main stages. The first involves discretizing the problem, while the second utilizes a numerical strategy to solve this system of discrete equations. Usually, no real interplay occurs between the discretization and solution processes. This results in enormous waste; the discretization process is unable to predict the required resolution and, consequently, the resulting algebraic system is not ‘optimal’. It is often unnecessarily large, while solution accuracy usually remains rather




Figure 4.2: A piece of non-uniform grid (a) and the uniform levels it is made of (b, c, d, e, f). In essence, a nonuniform grid is a union of uniform subgrids. However, unlike the traditional grids utilized in a multigrid, the subgrids do not necessarily extended over the same domain.

low. The ‘adaptive multigrid’ however, allows the interplay between the discretization process and the solution strategy, thereby making both more effective. The procedure can be considered a multigrid process in which the fine grid covers only part of the domain. It is the limited extent of the fine grid that provides the benefits to the method.

Non-uniformity Generated by Uniform Grids

Non-uniform resolution is needed in many, perhaps most, practical problems. Grid refinement is often required near singular points, in boundary layers and so on. Coarse grids can be used where the solution is smooth, or in regions away from locations where the solution is accurately needed. In this study, we recover the flexibility of non-uniform grids by generating a composite grid made up of a collection of patches, with each patch being a uniform grid. The domain covered by any patch may be only a small subdomain of the coarser grids. In such a structure, the effective mesh-size in each neighborhood will be that of the finest grid covering it.



To illustrate this concept, consider a simple domain, discretized by quadrilateral elements, made up of five grid levels. Suppose grid level one and two extend over the whole domain, as is shown in Figure 4.2b/c. This is common practice in standard multigrid programs. The grid at level three however does not span the entire domain (Figure 4.2d), being confined to the vicinity of the reentrant corner, where higher resolution is required to maintain solution accuracy. Grids four and five are also restricted to this region of the domain, becoming more limited in their extent as resolution becomes finer (Figure 4.2e/f). Thus, the final non-uniform grid, displayed in Figure 4.2a, is made up of five distinctive grid levels.

This structure is highly flexible, since local grid refinement (or coarsening) is done by extending (or contracting) uniform grids, which is relatively easy and inexpensive to implement. Furthermore, the use of partial grids leads to a considerable saving in both computational memory and operations, especially when only a small portion of the domain needs to be refined. There does, however, appear to be a certain waste in the proposed system, as one function value may be stored several times, when its grid point belongs to several levels. This is not the case. Firstly, the amount of such extra storage is small, being less than 2^{-d} of the total storage, for a d -dimensional problem (Brandt 1977). Moreover, the stored values are exactly those needed for the multigrid solution process. The major benefit of such an approach, therefore, is that it allows full integration of the grid adaptation process into the fast multigrid solver (using its levels also as in a multigrid solver), whereas other approaches rely on switching back and forth between solving and discretizing (hence their total work is not proportional, in principle, to the number of points of the final grid). Furthermore, by utilizing uniform grids, recurrent operators can be used for both relaxation and transfer procedures and a simple data structure can be employed. It should be noted that, in addition to being suitable for finite element schemes, the approach is also appropriate for finite difference and finite volume methods.

4.3.4 Implementation and Solution Procedures

The issues involved in approximating the governing equations on non-uniform grids are now considered. A 2-D domain is utilized to illustrate this process, discretized by quadrilateral elements, while a standard bisection refinement rule is employed, i.e. each quadrilateral element is split into four elements at the next grid level. The key ideas are presented in two-dimensions for ease of illustration and description. They are described further, in the framework of a 3-D spherical mantle dynamics code, in the following section.

A multigrid algorithm, based upon a hierarchy of non-uniform grids, must address three numerical issues:

1. The presence of irregular points, or ‘hanging nodes’, at the interface linking grids of different resolution.
2. Solution continuity during the prolongation and restriction inter-grid transfers.
3. The refinement strategy, i.e. will refinement (or coarsening) be performed adaptively, based upon *a posteriori* error indication criterion, or *a priori*, based on previous knowledge of the problem under investigation / the specific requirements of a simulation?

The first two points are addressed further in the following sections. The third point is dependent upon the user-specific requirements of an investigation and, consequently, need not be discussed in more detail.

Dealing With Hanging Nodes

A typical interface between grids of different resolutions is displayed in Figure 4.3. This could illustrate, for example, the interface between a level two grid and a level three grid, as would occur between the grids displayed in Figure 4.2c/d. There are two coarse-grid elements, but the right-hand element is subdivided into four fine-grid elements. The location




Figure 4.3: (a) The problem: One coarse-grid element interfacing with two fine-grid elements. The location of genuine unknowns is shown for both fine and coarse grids. (b) The problem, modified to show the dummy unknowns (virtual nodes) introduced for computational convenience. The variables shown here are not unknowns in our system of equations, but dummy variables that are introduced to allow consistent solution derivation at all genuine nodes.

of all genuine variables is shown in Figure 4.3a. By genuine we mean here unknowns that are associated with an approximation to some differential equation.

The figure illustrates that at the interface, the usual five-point stencils should be modified to involve both fine and coarse grid components. However, this is rather inconvenient. In order to remedy this situation a border of virtual nodes is generated in the coarse-grid region (see Figure 4.3b), as described by Thompson et al. (1992). These are introduced for computational convenience to regularize the structure of the unknowns.

The dummy unknowns are not updated during the relaxation process. These values are determined by interpolation only, utilizing standard inter-grid transfer operators. They are set at the beginning of the relaxation phase and act as boundary values, while the genuine variables, which are now separated from the interface, are updated.

Ensuring Solution Continuity During Inter-Grid Transfers

A second layer of virtual nodes is introduced around the first layer to ensure that mass is conserved during the inter-grid transfer processes of restriction. A nodal classification is introduced at this stage, to clarify the precise meaning and purpose of each node (see Figure 4.4).

The fine grid contains all nodes appertaining to the fine regions of the domain - the F (or fine) nodes in Figure 4.4. In addition, a boundary band consisting of two layers of nodes is included, which collects all nodes connected to the F -nodes through the relaxation (R -nodes) or transfer operators (T -nodes). The relaxation process is performed on the F -nodes and also involves R -nodes, however, as mentioned in the previous section, R -nodes are not updated during the relaxation process. The interpolation and restriction processes involve F -, R - and T -nodes. At this stage, the reader should focus on the coarse-grid nodes, encircled in red in Figure 4.4. To ensure the solution is accurately transferred to these nodes during restriction, solutions must be correct at the nodes involved in their 5-point stencil. Consequently, a solution must be generated at the so-called T -nodes. The transfer between grids can then be achieved via standard transfer operators and R - and T -nodes are initialized by direct interpolation from the coarse solution (C -nodes). These features allow inter-grid continuity conditions to be satisfied.

It is worth mentioning that an identical approach is employed by Lopez et al. (1998) and that a similar approach is also followed in Rivara (1984). However, in Rivara (1984), the T -nodes are not explicitly introduced and the values of R -nodes are computed directly from the coarse solution. In this way, a little storage is saved, but at the cost of either computational time or coding.




Figure 4.4: A subsection illustrating the nodal classification of the $h-1$ and h grids in our rectangular domain.

4.4 Implementation within TERRA

Having introduced the fundamentals of the adaptive multigrid in 2-D geometry, the implementation of a comparable scheme within a pre-existing 3-D spherical code is now described. Each step involved in this implementation is covered. However, before proceeding, the true nature of the executed scheme must be clarified.

As stated in the introduction, the main aim of this study is to demonstrate that multi-resolution techniques can be incorporated within pre-existing 3-D spherical mantle dynamics codes. This goal is achieved by utilizing the multigrid strategies presented above, however, two simplifications are incorporated. Firstly, the depicted implementation is not truly adaptive; fine and coarse grid regions are selected *a priori*. As a result, validating the code is a simpler task, while *a posteriori* error indication criteria, to determine where refinement is necessary, could be included at a later date. Secondly, refinement is restricted to the radial component of the spherical coordinate system. In essence, we implement a one-dimensional equivalent to the strategies introduced in the previous sections, albeit within a 3-D code. From this point on, therefore, the implemented method is termed 'radial

refinement'. It is important to note that the radial refinement strategies are equally valid for lateral refinement. An extension of this work at a later date is therefore possible.

Our code of choice is TERRA, a well established finite element solver that was first developed by Baumgardner (1983). A summary of the code is provided in Appendix E, while further details can be found in (Baumgardner 1983, 1985, Bunge & Baumgardner 1995, Yang 1997). TERRA solves the conservation equations of linear momentum (4.16), mass (4.17) and energy (4.18), inside a spherical shell with appropriate boundary conditions:

$$\nabla \cdot \tau - \nabla p + \rho g = 0 \quad (4.16)$$

$$\frac{\partial \rho}{\partial t} = -\nabla \cdot (\rho \mathbf{u}) \quad (4.17)$$

$$\frac{\partial T}{\partial t} = -\nabla \cdot (T \mathbf{u}) - (\gamma - 1) T \nabla \cdot \mathbf{u} + \frac{[\tau : \nabla \mathbf{u} + \nabla \cdot (k \nabla T) + H]}{\rho c_V} \quad (4.18)$$

where

$$\tau = \mu \left[\nabla \mathbf{u} + (\nabla \mathbf{u})^T - \frac{2I(\nabla \cdot \mathbf{u})}{3} \right] \quad (4.19)$$

and

$$p = p(\rho, T) \quad (4.20)$$

In the above equations p denotes pressure, ρ density, g gravitational acceleration, τ deviatoric stress, $\mathbf{u}$ fluid velocity, T absolute temperature, γ the Grüneisen parameter, k thermal conductivity, H volumetric radiogenic heat production and c_V specific heat at constant volume. μ represents the dynamic viscosity, which can be a function of temperature and pressure. Equation (4.19) is the linear constitutive law, which relates the stress field to the velocity, while Equation (4.20) represents the equation of state as a suitable function of density and temperature.

The spherical shell is discretized by an icosahedral grid. By projecting the regular icosahedron onto a sphere, the spherical surface can be divided into twenty identical spherical

triangles, or ten identical diamonds, each of which contains one of the ten triangles surrounding each of the two poles. Each triangle can subsequently be subdivided into four triangles by construction of great circle arcs between triangle side mid-points. This refinement process can be repeated, to yield an almost uniform triangulation of the sphere at any desired resolution. The grid is extended radially by placing several of these spherical shells above one and other, generating a mesh of triangular prisms with spherical ends.

TERRA uses the Galerkin finite element formulation, where lateral shape functions are represented by spherical barycentric coordinates. These are extended in the radial direction as simple one-dimensional linear shape functions. An Uzawa type pressure correction approach is coupled with a conjugate gradient algorithm to solve Equations (4.16) and (4.17), in terms of the aforementioned shape functions. The basis of this approach is that the velocity and pressure determined by solving Equation (4.16) alone should be corrected until Equation (4.17) is satisfied. The algorithm was originally proposed by Verfürth (1984) and is outlined in detail by both Atanga & Silvester (1992) and Ramage & Wathen (1994). It is computationally economical, exploiting a highly-efficient multigrid inner-solver, which takes advantage of the nested character of successive refinements of the icosahedral grid. The energy equation is solved by means of a flux-form finite difference method, in conjunction with the finite volume method, while time-stepping is accomplished via a second-order Runge-Kutta scheme. The code is written in the FORTRAN programming language and it has been parallelized using MPI (Bunge & Baumgardner 1995).

The fact that TERRA is built around a multigrid solver makes it an ideal candidate for investigating and validating the multi-resolution concepts introduced above. However, the executed techniques are equally valid for other 3-D spherical codes, providing that they are centered upon a multigrid solver. Consequently, although this study focusses on TERRA,

the findings will be of benefit to the wider geodynamical community.

4.4.1 Model Discretization

At this point a simple non-uniform discretization is presented (Figure 4.5a), which is later used to illustrate the key concepts of the radial refinement technique. It is important to note that the implemented method can deal with other discretizations, as illustrated by Figure 4.5b/c. However, the methods presented in this chapter focus on that shown in Figure 4.5a.

The spherical shell is separated into two distinctive regions; a high resolution upper mantle with coarser resolution in the lower mantle. Such a configuration minimizes bookkeeping, whilst allowing the multi-level process to be illustrated through simple one-dimensional diagrams. In spite of its simplicity, this non-uniform discretization yields significant benefits over the conventional quasi-uniform structure. The original icosahedral discretization leads to considerable simplifications in both storage, data structure, parallelization and solution strategies. However, the scheme also has some disadvantages:

1. Grid resolution can only be changed in fixed step sizes, globally. For example, increasing grid resolution by one stage results in an increase in the number of nodes by a factor of eight. Eight times more RAM and approximately sixteen times more processing time is therefore needed to solve problems at the next level of refinement, the increased factor in CPU-time resulting from the need to decrease the time-step. A local increase or decrease in resolution cannot be accommodated.
2. Element sizes and inter-nodal distances are dependent upon radius, with grid points denser at the inner surface of the spherical shell than on the outer surface. A caveat to this is that, as a consequence of the Courant-Friedrichs-Levy stability criterion, the length of a dynamically controlled time-step is dictated by the smaller elements at the base of the shell. This is far from ideal.

The non-uniform discretization illustrated in Figure 4.5a goes some way towards resolving these shortcomings. To begin with, resolution is enhanced in the upper mantle. This




Figure 4.5: (a) An illustration of the grid configuration implemented in this study; high resolution upper mantle, with coarser resolution in the lower mantle. The final solution is derived from distinctive local grids. (b) and (c) are examples of alternative grid configurations that could be accommodated within the present setup.

counteracts the setup of the original grid, which is characterized by superior resolution in the lower mantle. As a consequence, the outer shell boundary can be resolved to the same, if not a greater extent than the inner shell boundary. Secondly, element size and inter-nodal distances are more consistent over the entire domain, meaning that the dynamically controlled time-step is not necessarily dictated by smaller elements at the base of the shell. Finally, and perhaps most importantly, since resolution can be increased locally, albeit, in the radial direction, computational demands at the next level of refinement are significantly reduced. This is of great benefit, allowing one to simulate more demanding problems than were previously possible.

There are also significant benefits to this non-uniform discretization from a geodynamical perspective. The viscosity of Earth's mantle is believed to increase significantly with depth (e.g. Mitrovica & Forte 2004). As a consequence, one would expect convection in the upper

mantle to be characterized by fine-scale activity, with longer wavelength features becoming more prevalent with depth. The discretization presented, therefore, with fine resolution in the upper mantle, fits in perfectly with such behavior.

A further advantage to this configuration is that it fits in perfectly with the current setup of TERRA. By utilizing uniform grids at each level of refinement, standard operators can be used for both relaxation and inter-grid transfer procedures. Consequently, no major modifications are needed to the code. Furthermore, the proposed discretization is ideal for the parallelization and domain decomposition inherent to TERRA, retaining the equal load-balance of the original scheme. This is of utmost importance for computational efficiency.

In conclusion, the model non-uniform discretization presented not only allows the fundamental methodological concepts to be introduced in a straightforward manner, but it yields significant computational benefits and also advantages for geodynamical investigations.

4.4.2 Converting TERRA's Multigrid From CS To FAS

The original version of TERRA employs a correction scheme multigrid algorithm to solve an algebraic equation like (4.1). As is clear from Section 4.3.3, this scheme must be modified to the FAS mode of multigrid processing before the radial refinement strategies can be incorporated, since in parts of the domain not covered by the finest grid, the coarser grid must show the full solution, not just a correction. The revised algorithm, as well as the original, are presented below. For simplicity, this discussion is restricted to a two level system, consisting of a fine grid (grid spacing h) and one coarse grid (grid spacing $H = 2h$). Let us begin with the original CS scheme. Suppose we are given an estimate, u_0 , for the solution of Equation (4.1). The residual is then given by:

$$r^h = f^h - A^h u_0 \quad (4.21)$$

The solution to the equation:

$$A^h e^h = r^h \quad (4.22)$$

is sought, where e^h represents the error associated with the approximate solution - the correction that should be added to u_0 to improve the estimate for u . A crucial aspect of the multigrid method is to find a coarser grid correction field e^H from:

$$A^H e^H = r^H \quad (4.23)$$

where A^H and r^H are the forward operator and the residual vector on the coarser grid. Once we have A^H and r^H , we can determine e^H either through a relaxation method, or, when the matrix is manageable, solving the equation exactly. A *pseudo-inverse* method is utilized in TERRA at this stage (see Baumgardner 1983, for further details).

The fine grid representation $e^{h'}$ of e^H can be obtained by utilizing an interpolation operator. To improve the correction field, $e^{h'}$, thus obtained, a second order residual is formed on the fine-grid as:

$$r^{h'} = r^h - A^h e^{h'} \quad (4.24)$$

A second order correction term $e^{h''}$ is then found by relaxing the equation:

$$A^h e^{h''} = r^{h'} \quad (4.25)$$

The final correction field on the fine grid is given by:

$$e^h = e^{h'} + e^{h''} \quad (4.26)$$

One multigrid iteration is completed when this term is added to u_0 . The corrected solution is fed back into Equation (4.21) and the iterating process is continued until the residual becomes acceptably small. This is an example of a *V-cycle*.

As was noted earlier, for linear problems, the CS and FAS modes are fully equivalent. Since TERRA's forward operator is indeed linear, conversion of the multigrid to the FAS mode is relatively straightforward. The only difference between the schemes is that, instead of e^H , the FAS mode is performed in terms of another coarse-grid variable, u^H , and the full solution is stored at all grid levels. Switching to the FAS scheme is therefore simply a case of storing the approximation, u , at each grid level and adding to it the correction field at that level. The full algorithm is presented below, again for a two level system.

The first stage involves simple injection of the fine grid solution to the coarse grid. Assuming that we are given an estimate, u_0 , for the solution of Equation (4.1), we form u^H , the coarse-grid solution, as:

$$u^H = I_h^H u^h \quad (4.27)$$

where I_h^H is a fine-to-coarse injection operator. The next stages are identical to those of the original CS scheme. Equations (4.21)–(4.23) are solved to obtain the coarse grid correction, e^H . This coarse grid correction e^H , is then added to the coarse grid solution array u^H as:

$$u^H \leftarrow u^H + e^H \quad (4.28)$$

In this way, the full solution is stored on the coarser grid, as opposed to the correction alone. Since the FAS interpolation back to the fine grid simplifies to that of the CS for linear operators, e^H need not be modified before interpolation to the fine grid. The remaining stages of the implemented FAS scheme are then identical to the CS, with Equations (4.24)–(4.26) solved to obtain an improved fine grid correction field, which is then added to the fine grid solution, u^h . The final solution is identical to that obtained in the CS scheme.

4.4.3 Adaptive Solution Strategy for Force-Balance Equation

The strategy employed to solve the force-balance equation strongly resembles the Uzawa conjugate gradient algorithm utilized in the original version of TERRA, with the minor difference that fine-grid calculations are performed on local grids as opposed to global grids. These local calculations are executed using the FAS multigrid scheme presented in the previous section, in conjunction a series of virtual nodes and a radial refinement technique, which is essentially a one-dimensional equivalent to the adaptive multigrid concepts introduced in Section 4.3.3. The fundamental components to the radial refinement approach are presented here, employing a nodal classification similar to that presented in Section 4.3.4.

A radial section through the spherical shell is displayed in Figure 4.6a, illustrating the radial location of all genuine nodes i.e. those radial layers represented in the final solution. The final grid is non-uniform, containing both fine (grid-level h) and coarse (grid level $h - 1$) components, upper mantle solutions obtained at grid level h (red shaded nodes in Figure 4.6b) and lower mantle solutions calculated at grid level $h - 1$ (red shaded nodes in Figure 4.6c). The FAS multigrid scheme yields the approximation at all nodes, be it coarse or fine, providing a gateway to non-uniform resolution. However, the conventional FAS algorithm must be coupled with a suitable boundary interpolation strategy to make such a non-uniform scheme possible. This modified FAS multigrid algorithm is given below, for a three-level system, i.e. grids h , $h - 1$ and $h - 2$ in Figure 4.6. Given an initial estimate of the solution, u_0 , at all genuine nodes:

1. Locally inject solution to grid $h - 1$ by Equation (4.27), where grids h and $h - 1$ overlap. Next, inject solution from grid $h - 1$ to grid $h - 2$. The solution now exists at all grid levels.
2. Compute fine grid residual from Equation (4.21). A conventional scheme would utilize all nodes displayed in Figure 4.6b at this stage. However, since N -nodes do not exist in




Figure 4.6: (a) A radial section, drawn along AB, illustrating the final non-uniform solution grid in our simulations. (b) Radial configuration of grid level h , a local grid that spans the upper half of the mantle. Nodes shaded in red are those intrinsic to the final solution - *F*-nodes. Nodes utilized as boundary values during local relaxation are labeled *R*-nodes, while those employed to ensure solution continuity during inter-grid transfer are labeled *T*-nodes. *N*-nodes are nodes that would occur in the regular multigrid formulation, but do not in the adaptive multigrid formulation. In other words, they do not actually exist and are included for illustrative purposes only. (c) Radial configuration of grid level $h-1$, i.e. coarse grid *C*-nodes. This is a global grid, covering the whole mantle. Nodes shaded in red are part of the final solution, whilst nodes shaded in grey are utilized during the multigrid process but do not explicitly contribute to the final solution. (d) Radial configuration of grid level $h-2$, another global grid which is involved in the multigrid scheme but does not explicitly contribute to the final solution.

the modified formulation this cannot be done and a local residual calculation is needed. A virtual *R*-node is introduced via interpolation, directly beneath the lowermost *F*-node. By acting as a boundary value, residuals can be correctly calculated at all *F*-nodes. The next stage would be to locally project this residual to grid level $h-1$. However, at present, the residual cannot be accurately transferred to the lowermost grey *C*-node at this level. The scheme must therefore be modified. To guarantee accurate residual projection, a second virtual node is introduced, the *T*-node. This ensures that residuals are accurately calculated at *R*-nodes, since the *T*-nodes act as boundary values during fine grid residual calculation. Such a scheme ensures solution continuity during projection.

3. Locally project residual at grid level h to grid level $h-1$. In regions where grids are

not nested/do not overlap, calculate coarse-grid residual directly.

4. Project full coarse-grid residual to grid level $h - 2$.
5. Calculate coarse grid correction by solving Equation (4.23).
6. Add correction to current approximation at grid level $h - 2$ via Equation (4.28).
7. Interpolate correction to grid level $h - 1$ and improve by solving Equations (4.24)–(4.26).
8. Add improved correction field to current approximation at grid level $h - 1$ via Equation (4.28).
9. Locally interpolate correction field to F - and R - nodes at grid level h . Improve the correction by solving Equations (4.24)–(4.26) locally, utilizing R -nodes as boundary values during the solution process.
10. Add improved correction field to current approximation via Equation (4.28)
11. Improve solution at fine grid F -nodes via relaxation. R -nodal values are set at the beginning of the relaxation phase and act as boundary values, while genuine variables (F -nodes), now separated from the $h/h - 1$ interface, are updated.
12. The corrected solution is fed back into Equation (4.27) and the iterating process is continued until the residual becomes acceptably small.

In summary, fine grid nodes (i.e. those in the upper mantle) are represented by F -nodes. Nodes utilized as boundary values during local relaxation are labeled R -nodes, while those employed to ensure solution continuity during inter-grid transfer are labeled T -nodes. C -nodes represent coarse grid nodes (i.e. those in the lower mantle), some of which directly contribute to the final solution.

The localized Uzawa solution algorithm has not been fully presented here. To avoid repetition, we have focussed on the multigrid inner-solver. The full Uzawa scheme is identical to that originally implemented in TERRA, excluding the localization of certain calculations

and minor interpolation stages. *R*- and *T*-nodes are employed to ensure solution accuracy during this process, exactly as described in the above multigrid algorithm. Consequently, it is not necessary to describe this process repetitively for each step of the Uzawa solution procedure.

4.4.4 The Multi-Level Energy Equation

The methods used to treat the advection, thermal conduction and time integration of the energy conservation equation over grids of variable resolution are next described. The methods used are essentially identical to those utilized in the original version of TERRA (see Appendix E.7). This section however will outline the subtle modifications made to the original solution strategy to account for the irregular grid.

Figure 4.7 illustrates the nodal classification employed whilst solving the multi-level energy equation. *S*-nodes are those nodes represented in the final solution. Note once again that the solution in the upper half of the domain is derived from the fine grid level h while the solution in the lower half is calculated at the next grid level $h - 1$. *G*-nodes are those where all components involved in determining the advective and conductive fluxes at the *S*-nodes need to be accurately prescribed. These include nodal velocities, temperatures, and temperature gradients. Clearly, in order to determine the true temperature gradient at the *G*-nodes, temperatures must be correctly specified at the connected nodes - these are the *P*-nodes. *I*-nodes are coarse grid (level $h - 1$) nodes not directly involved in the energy equation solution procedure. As in Figure 4.6a, *N*-nodes are the nodes that would occur in the regular formulation, but do not appear in the modified formulation.

The modified solution algorithm for the advective heat flux, based around the aforementioned nodal classification is summarized below (assuming a prescribed / pre-existing ther-




Figure 4.7: Nodal classification employed whilst solving the multi-level energy equation. *S*-nodes are those implicated in the final solution, whilst *G*- and *P*- nodes are ‘virtual’ nodes, which, in essence, only exist to ensure that the solution at the *S*-nodes is accurate. *I*-nodes are coarse grid nodes not directly involved in the energy equation solution procedure. Note that the *N*-nodes do not actually exist. They are included for illustrative purposes only - they would occur in a regular multigrid formulation, but they do not occur in the adaptive formulation.

mal field at all *S*-nodes):

1. Generate *G*- and *P*-node temperatures at grid level $h - 1$ via simple injection from corresponding fine grid nodes.
2. Generate the temperature at the fine grid *P*-node through injection from corresponding coarse-grid node. Utilize value at this *P*-node together with that at the fine grids lowermost *S*-node, the node at grid level h contained within a red box, to derive a temperature value at the fine grid *G*-node, via linear interpolation.
3. Loop over all fine-grid *S*-nodes, from the surface to the base of the fine grid region, calculating the rate of heat advection for all elements. To ensure accuracy of the advective heat flux at the lowermost *S*-node, the density, specific heat, temperature and velocities must be correctly prescribed at the associated *G*-node. Temperature gradients must also be accurate at the *G*-node.
4. Loop over all coarse-grid *S*-nodes, beginning at the uppermost node (the node at grid level $h - 1$ contained within a red box), and ending at the lowermost node, calculating the advective heat flux. Again, the accuracy of the calculation at the uppermost *S*-

node can only be guaranteed when the density, specific heat, temperature, temperature gradient and velocities are correctly prescribed at the associated G -node. As with the fine grid, P -nodes are only involved at this stage of the solution process, ensuring that temperature gradients calculated at G -nodes are accurate.

This nodal classification and solution algorithm ensures that continuity between fine and coarse grids is satisfied, since the flux calculated at each S -node is accurate. It should be noted that the methods used to treat thermal conduction are analogous to those utilized in tracking thermal advection.

Time Integration

An explicit Runge-Kutta time stepping scheme has been adopted in this study. Such schemes can be recast into an m -stage scheme of the form:

$$\Delta T^{n+1} = \alpha_i \Delta t r (T^n + \Delta T^{n+i-1}) \quad i = 1, m \quad (4.29)$$

where T is the temperature, r represents the residual in the energy equation and Δt is the time step. A fourth order scheme ($m = 4$) is utilized here, providing fourth-order accuracy and requiring derivatives at four points in each time interval. This contrasts with the original second-order scheme employed by TERRA, which provides second-order accuracy, whilst requiring derivatives at two points in each interval.

The coefficients, α , are selected to produce a scheme that is suitable for multigrid (Lohner 2001), such that $\alpha_1 = 0.25$, $\alpha_2 = 0.5$, $\alpha_3 = 0.5$ and $\alpha_4 = 1$. Written out in full, the time-stepping algorithm is:

$$\begin{aligned} \Delta T^n &= 0 \\ \Delta T^{n+1} &= \alpha_1 \Delta t r (T^n + \Delta T^n) \\ \Delta T^{n+2} &= \alpha_2 \Delta t r (T^n + \Delta T^{n+1}) \end{aligned}$$

$$\begin{aligned}
 \Delta T^{n+3} &= \alpha_3 \Delta t r (T^n + \Delta T^{n+2}) \\
 \Delta T^{n+4} &= \alpha_4 \Delta t r (T^n + \Delta T^{n+3}) \\
 T^{n+1} &= T^n + \Delta T^{n+4}
 \end{aligned} \tag{4.30}$$

where r is the residual of the energy equation.

The radial refinement strategies, nodal classification and solution algorithms described in the previous sections allow the incorporation of multi-resolution within TERRA. Furthermore, the nature of their formulation ensures that no major modifications need to be made to the original arrangement of TERRA and its nodal stencils. This is of great importance, illustrating that local regions of high resolution can be incorporated into pre-existing solvers without the need to reformulate and restructure large amounts of code. The validation of these methods is covered in the following section.

4.5 Validation of TERRA and Radial Refinement Scheme

The accuracy of TERRA and the radial refinement algorithm is now demonstrated by comparing results from the code with previously published analytical and numerical solutions for two different problems. The first problem involves a simple comparison of stationary convection at relatively low Rayleigh numbers, in isoviscous media (e.g. Bercovici et al. 1989, Ratcliff et al. 1996, Choblet et al. 2007). The second problem is for Stokes flow with both an isoviscous and layered viscosity structure, for which quasi-analytic solutions can be obtained via propagator matrix methods (Hager & O'Connell 1981, Richards & Hager 1984).

4.5.1 Thermal Amplitude Tests

Two classes of 3-D flows, that exhibit polygonal symmetry, are investigated, predicted analytically by Busse & Riahi (1982) and confirmed numerically by Bercovici et al. (1989). These symmetric flows have been termed tetrahedral and cubic on the basis of the relative positions of the upwelling plumes in the flow. Tetrahedral patterns, which exhibit a degree $l = 3$ and order $m = 2$ dominant spherical harmonic signature, have four upwelling plumes that intersect the upper surface of the spherical shell at points corresponding to the vertices of a tetrahedron enclosed by the outer spherical surface. The six upwelling plumes of the cubic planform intersect the outer surface at points that correspond to the centers of the faces of a cube enclosing the spherical shell, creating a flow pattern with two dominant spherical harmonic modes: $l = 4, m = 0$ and $l = 4, m = 4$. Numerical studies (e.g. Bercovici et al. 1989, Ratcliff et al. 1996, Harder & Christensen 1996, Zhong et al. 2000, Yoshida & Kageyama 2004, Stemmer et al. 2006, Choblet et al. 2007) have found cubic and tetrahedral flows to be steady over a large range of Rayleigh numbers.

In order to test that TERRA, and subsequently our modified solution algorithm, produces accurate results we reproduce these two cases. The first test case is for tetrahedral symmetry at $Ra = 7,000$, while the second is for cubic symmetry at $Ra = 3,500$. A comparison of Nusselt numbers (Nu) and non-dimensional RMS velocities (V_{RMS}) are made. Nusselt numbers are determined by solving the time-dependent energy equation until the relative variation in the Nusselt number between two consecutive time-steps is less than 10^{-6} . RMS velocities are calculated once Nusselt numbers are deemed to have converged.

Results are presented in Table 4.1. An excellent agreement is observed between our results and those obtained in previous studies, suggesting that both the code and radial refinement strategies are valid and accurate. A comforting observation is the small difference between

(a)			
Reference	Nu_{upper}	Nu_{lower}	V_{RMS}
B89	3.4657	3.5293	-
H96	3.4957	-	-
R96	3.4423	-	32.19
Z00	3.519	3.472	-
Y04	3.443	-	32.05
S06	3.4864	3.4864	32.59
C07	3.4814	3.4717	32.74
D08	3.5166	3.4833	32.35

(b)			
Reference	Nu_{upper}	Nu_{lower}	V_{RMS}
B89	2.7954	-	-
H96	-	-	-
R96	2.8306	-	18.86
Z00	-	-	-
Y04	2.8830	-	18.48
S06	-	-	-
C07	2.8640	2.8948	19.55
D08	2.8135	2.7864	18.49

Table 4.1: Thermal amplitude benchmark comparisons for an isoviscous fluid: (a) Case 1, tetrahedral symmetry at $Ra = 7 \times 10^3$, (b) Case 2, cubic symmetry at $Ra = 3.5 \times 10^3$. Nu_{upper} and Nu_{lower} represent the upper and lower boundary Nusselt numbers, respectively, while V_{RMS} denotes the non-dimensionalized RMS velocity. Abbreviations in the first column refer to the studies used for comparison: B89 (Bercovici et al. 1989), H96 (Harder & Christensen 1996), R96 (Ratcliff et al. 1996), Z00 (Zhong et al. 2000), Y04 (Yoshida & Kageyama 2004), S06 (Stemmer et al. 2006), C07 (Choblet et al. 2007). D08 is the present study.

inner and outer shell surface heat fluxes, indicating that the scheme is globally conservative.

These tests suggest that the methodology is valid, however, global averages are not a sufficient criterion for a precise benchmark. Further validation is necessary with a more robust benchmark presented in the following section.

4.5.2 Stokes Flow

The second benchmark problem exclusively tests the Stokes flow, in which Equations (4.16) and (4.17) are solved, for both isoviscous and layered-viscosity structures. Comparisons are made with quasi-analytical solutions to these equations, derived via propagator matrix methods (e.g. Hager & O'Connell 1981, Richards & Hager 1984). We specifically examine normalized poloidal velocity coefficients obtained in response to a spherical harmonic

μ_1/μ_2	(l, m)	V_{oa}	V_{om}	V_{ia}	V_{im}	$Error_o(\%)$	$Error_i(\%)$
1	(2.0)	6.028E-10	6.010E-10	-7.112E-10	-7.061E-10	0.30	0.72
	(4.0)	3.094E-10	3.121E-10	-3.388E-10	-3.360E-10	0.87	0.83
	(8.0)	6.844E-11	6.943E-11	-5.731E-11	-5.801E-11	1.45	1.22
	(16.0)	4.553E-12	4.673E-12	-1.899E-12	-1.844E-12	2.57	2.90
0.3	(2.0)	9.303E-10	9.281E-10	-8.368E-10	-8.305E-10	0.24	0.75
	(4.0)	5.646E-10	5.691E-10	-4.404E-10	-4.369E-10	0.80	0.79
	(8.0)	1.272E-10	1.302E-10	-6.943E-11	-7.013E-11	2.36	1.01
	(16.0)	7.269E-12	7.700E-10	-1.971E-12	-2.040E-12	5.93	3.50

Table 4.2: Surface normalized harmonic coefficients for poloidal velocity. μ_1/μ_2 denotes the ratio between upper (μ_1) and lower (μ_2) mantle viscosities, while (l, m) is the spherical harmonic function, of degree l and order m . V_{oa} and V_{om} represent analytical and model solutions at the outer boundary, V_{ia} and V_{im} the equivalent solutions at the inner boundary, whilst $Error_o$ and $Error_i$ represent the percentage error of the model solution at the outer and inner boundaries respectively.

temperature perturbation at a specified depth in the spherical shell.

The problem is set up as follows. The inner radius is set to mimic that of Earth's core mantle boundary, $r_i = 3480\text{km}$, while the outer radius is set to equal that of Earth's surface, $r_o = 6370\text{km}$. Free-slip mechanical boundary conditions are specified at both surfaces. The driving force is a delta function temperature perturbation in radius, defined as:

$$T = \delta\left(r - \frac{r_i + r_o}{2}\right)Y_{lm}(\theta, \phi) \quad (4.31)$$

Here, Y_{lm} is the spherical harmonic function of degree l and order m . Two cases are investigated, at a range of spherical harmonic degrees. The first case simulates a uniform viscosity structure (1×10^{22} Pa s) throughout the shell (i.e. $\mu_1/\mu_2 = 1$, where μ_1 and μ_2 are upper and lower mantle viscosities, respectively) while the second case models a viscosity structure that consists of two distinct layers, with μ_1 having a relative viscosity of $0.3 \times \mu_2$ ($\mu_2 = 1 \times 10^{22}$ Pa s) and spanning the upper quarter of the spherical shell.

Results are presented in Table 4.2, demonstrating that for Stokes flow with both isoviscous and layered viscosity structures, normalized poloidal velocity coefficients agree well with analytical solutions. As would be expected from theory, the agreement diminishes as one goes to higher and higher harmonic degrees. However, for simulations at harmonic degrees




Figure 4.8: Plots of solution error against grid resolution, for Stokes flow, at various spherical harmonic degrees (SHD), for (a) uniform viscosity structure and (b) depth-dependant viscosity structure. Blue represents SHD 16, Red SHD 8, Green SHD 4 and violet SHD 2. As expected, results converge towards the analytical solution as grid resolution increases.

2 and 4, results agree to within $< 1\%$, which is comparable to the findings of Zhong et al. (2000). Figure 4.8 also demonstrates that results are convergent, approaching the analytical solution as grid resolution increases. These results, along with those presented in Section 4.5.1, give us great faith that the code and hence the new techniques are valid.

4.6 Computational Efficiency

In this section, a particular case of 3-D mantle flow is examined to quantify the computational efficiency, in terms of both RAM and CPU-time, of the aforementioned radial

Resolution (mt)	Radial Layers (nr)	N° Nodes	Nu_{Top}	Nu_{Bot}	N/D RMS_{Total}	RMS_{Surf} (m/s)
64	32	1352076	8.67	7.78	190.45	6.37×10^{-11}
128	64	10650380	9.09	8.43	191.69	6.46×10^{-11}
256	128	84542988	9.22	8.56	191.76	6.49×10^{-11}
128/64	17/24	3768812	9.06	8.28	191.62	6.43×10^{-11}
256/128	33/48	29492172	9.20	8.47	191.71	6.48×10^{-11}

Table 4.3: Final Nusselt numbers and RMS velocities from a series of simulations at various grid resolutions. The upper section represents results obtained with a uniform discretization, while the lower section encompasses results attained with non-uniform resolution, where, for example, 128/64, has a resolution of $mt = 128$ in the upper 25% of the spherical shell and $mt = 64$ elsewhere. Nu_{Top} and Nu_{Bot} represent Nusselt numbers at upper and lower surfaces respectively, while N/D RMS_{Total} denotes the non-dimensional RMS velocity across the whole mantle and RMS_{Surf} denotes the surface RMS velocity.

refinement scheme. Results are compared with those obtained on regular, uniform discretizations. The example considered is a symmetric flow exhibiting a tetrahedral pattern at $Ra = 70.000$, which, excluding the Rayleigh number, is identical to Case 1 in Section 4.5.1. Calculations are performed at various resolutions, allowing simple comparisons between different discretizations. It should be noted that, for non-uniform cases, fine resolution is restricted to the upper 25% of the spherical shell, as illustrated in Figure 4.5b. Each simulation is terminated after 3×10^9 model years.

Prior to examining results, it is important to demonstrate that all calculations, regardless of the solution procedure and grid resolution, evolve towards the same end member. However, certain aspects of TERRA need to be clarified before continuing. Within the code, refinements to the grid are referenced by ‘the number of grid intervals along an icosahedral diamond edge’. This value is labeled mt , defining the number of data points present in a spherical surface and, hence, lateral resolution. The number of data points on a spherical surface is given by $10(mt^2 + 1) + 2$, since there are 10 icosahedral diamonds on each surface and two polar nodes. The grid is extended radially by placing several of these spherical shells above one and other, generating a mesh of triangular prisms with spherical ends. In this study, nr , the number of radial layers, is set to $mt/2$. The total number of nodes in the spherical shell is then given by $(nr + 1)[10(mt^2 + 1) + 2]$. Note that there are $nr + 1$

radial layers.

Table 4.3 displays the Nusselt numbers and *RMS* velocities obtained in a series of simulations at various grid resolutions. These values, although varying slightly with resolution as one would expect, demonstrate that calculations have evolved concurrently. This point is reinforced by examining the visual patterns displayed in Figure 4.9, where the temperature planforms exhibited by each simulation are consistent. These models are therefore suitable for comparisons of computational efficiency. It should be noted that calculations have only evolved for 3×10^9 years and have not yet reached a quasi steady state. This accounts for the fact that Nusselt numbers at the inner and outer boundaries are not comparable.

Let us begin with the results obtained from a uniform discretization allowing us to better understand the workings of TERRA (displayed in Table 4.4a and Figure 4.10). As grid resolution is increased from one level of refinement to the next, for example, from $mt = 64$ to $mt = 128$, the number of nodes increases by a factor of ≈ 8 (Figure 4.10a). In addition, inter-nodal spacing decreases by a factor of 2. As a consequence, the dynamically controlled time-step must decrease to ensure numerical stability, with around twice the number of time steps needed to reach a specified model time (Figure 4.10b). From these simple statements, one could estimate that a simulation will take ≈ 16 times longer at the next level of refinement. This, of course, would only be the case if TERRA scaled perfectly, which, as expected, is not the case. In moving from $mt = 64$ to $mt = 128$, and $mt = 128$ to $mt = 256$, total CPU-time (i.e. CPU-time per processor $\times$ N^o processors) increases by a factor of 16.76 and 21.27 respectively. This demonstrates an element of overhead in the communication, which increases with successive refinements (Figure 4.10c). However, since the overhead is minimal, this is not a major issue. A reassuring point to note from Figure 4.10d is that as one goes to higher resolution, the amount of work per node for each time step remains




Figure 4.9: Temperature snapshots at 3×10^9 years, from a series of simulations at various resolutions. The images show a radial surface 100km above the CMB, a cross-section and a hot yellow iso-surface at 2000K. Solid lines at the surface represent coastlines. (a) is the initial condition for all models. Note that, is not drawn to scale, the scale has been reduced to emphasize minor variations in temperature that would not otherwise be visible; (b) and (c) are the planforms yielded by radially refined simulations at $mt = 128/64$ and $mt = 256/128$ respectively; (d)-(f) are those yielded by uniformly discretized simulations at (d) $mt = 64$, (e) $mt = 128$ and (f) $mt = 256$. The figure demonstrates that each model has evolved in unison, displaying only slight variations with resolution.

virtually identical. This illustrates that the multigrid is attaining a rate of convergence that is independent of the number of grid points.

Turning our attention to the efficiency of the radially refined discretization and associated solution algorithms, we find a significant reduction in the time taken for an identical calculation with the modified formulation. From Table 4.4b we see that the calculation at $mt = 128/64$ takes $\approx 83,000$ seconds as opposed to $\approx 215,000$ on a uniformly discretized $mt = 128$ grid (Table 4.4a), while a calculation at $mt = 256/128$ takes $\approx 1.2 \times 10^6$ seconds versus $\approx 4.6 \times 10^6$ seconds at $mt = 256$. Therefore, total CPU-time decreases by a factor




Figure 4.10: Plots of various variables against grid resolution (mt) for uniformly discretized simulations in TERRA. (a) the number of nodes in the final mesh; (b) the number of time steps required to achieve a model time of 3×10^9 years; (c) total CPU-time required; (d) work per node at each time step.

of between 2 and 4, with the benefits increasing with resolution. In addition there is a significant reduction in the number of nodes and, hence, the amount of RAM needed to solve the problem. These results, of course, are only relevant to simulations that are valid within the modified framework, i.e. where the non-uniform grid is sufficiently fine globally to resolve the simulation's active features. Since the viscosity of Earth's mantle is believed to increase significantly with depth (e.g. Mitrovica & Forte 2004), one would expect upper mantle convection to be dominated by fine-scale activity, with longer wavelength features

(a)				
Resolution (<i>mt</i>)	N° Nodes	Total CPU-Time (s)	N° Time-Steps	Final Time-Step (yrs)
64	1352076	12872	948	2.90×10^6
128	10650380	215712	1953	1.35×10^6
256	84542988	4588800	4512	6.55×10^5

(b)				
Resolution (<i>mt</i>)	N° Nodes	Total CPU-Time (s)	N° Time-Steps	Final Time-Step (yrs)
128/64	3768812	83168	1295	2.02×10^6
256/128	29492172	1226624	2631	9.62×10^5

Table 4.4: The resolution, number of nodes, total CPU-time, number of time steps and final dynamically determined time step from (a) uniformly discretized simulations, and (b) multi-resolution simulations, after 3×10^9 years.

more prevalent in the lower mantle. The advocated scheme would therefore fit in perfectly with such behavior, yielding great benefits for geodynamical investigations. In addition, as mentioned previously, the modified discretization counteracts the setup of the original grid, which is characterized by finer resolution in the lower mantle.

In examining the influence of the modified solution algorithm on the work per node at each time step, we find that it is 30-70% less efficient than the original code, with these figures again improving with resolution (since the scheme will only be used for calculations at high resolution, this is a reassuring point). A certain overhead is expected, since the modified scheme involves interpolation of values to and from ghost nodes, localization routines etc... This result therefore is not surprising. Nonetheless, after combining all contributing factors, for a simple, first-order test of computational efficiency, the advocated scheme is proven to be significantly more efficient than the original. It allows for a larger stable time step in simulations, meaning that less time steps are required to achieve a specified model time. This is clearly the predominant controlling factor, since, regardless of the decreased efficiency in the work per node at each time step, the scheme uses significantly less CPU-time than the original.

In conclusion, the radial refinement strategy allows enhancement of upper mantle resolution,

in a computationally efficient manner. One can attain a solution to a non-uniform problem, consisting of a high resolution upper mantle with coarse resolution elsewhere, in little more time than a run with a globally uniform coarse grid. Due to its resourcefulness in terms of RAM, the scheme allows one to attempt calculations that would otherwise be impractical. One can therefore achieve higher upper mantle resolutions than would previously be the case. In addition, the reduced CPU-time needed to solve a given problem, when compared to the original scheme, allows one to solve problems quickly. Consequently, the technique will be ideal in investigations requiring parameter space mapping and/or numerous simulations.

4.7 Conclusions

A novel method for introducing non-uniform resolution to pre-existing 3-D spherical mantle convection codes has been explored here. The method, based around the application of a multigrid solver on non-uniform structured grids, yields similar advantages to those observed with unstructured grids, being computationally highly efficient. However, unlike unstructured techniques, the advocated method negates the need to reformulate large sections of code, whilst also being conceptually simple and straightforward to implement. The proposed methodology has been validated and an excellent agreement is observed with results from a wide range of other studies. Our methodology is accurate and robust.

This investigation focusses on the finite element formulation and the well known geodynamical code TERRA. However, the method is equally applicable to finite difference and finite volume schemes and could be implemented in other 3-D spherical mantle dynamics codes, which are centered upon a multigrid solver. The findings of this study should therefore be of benefit to the wider geodynamical community.

The main objective of incorporating multi-resolution techniques within a pre-existing 3-D

spherical mantle dynamics code has been achieved. However, as the title to this chapter suggests, this study has only gone some way towards achieving automatic, error-guided multi-resolution. We firmly believe that this should be the focus of future research. As has been demonstrated recently in 2-D (Davies et al. 2007), such a capability would yield highly efficient solutions to a wide variety of mantle convection problems, for which, at present, results are unattainable due to computational limitations.

Chapter 5

Mobile and Steady Plumes Coexist at Earth-like Vigor

5.1 Abstract

Using a highly efficient, benchmarked, multi-resolution version of the 3-D spherical mantle dynamics code TERRA, we present the highest vigor global mantle convection simulations to date. We examine upwelling structures, focussing upon the nature of deep mantle plumes. Previous studies have shown long-lived, anchored, coherent upwelling plumes to be a feature of low to moderate vigor convection. Since more vigorous convection traditionally shows greater time-dependence, the fixity of upwellings would not logically be expected for non-layered convection at higher vigors. Surprisingly, such configurations have recently been observed, in models incorporating a layered viscosity structure. With upwelling dynamics believed by many to be closely linked to hot-spot magmatism, it is of great importance to ascertain whether or not these conclusions are valid at the dynamical regime of Earth's mantle, which convects at a Rayleigh number of order 10^9 . By utilizing computer power

efficiently, the modified code allows the simulations presented herein to replicate such vigor. Steady, long-lived plumes are a feature of such systems. However, mobile and ephemeral plumes are also observed. Such variable plume characteristics are consistent with hot-spot observations on Earth. Those plumes that do remain steady vary in intensity throughout the simulation. Such behavior is caused by an irregular supply of cold material to the core-mantle boundary region. Since subduction on Earth is a time-dependent and asymmetric process, our models suggest that subducting slabs might be partially responsible for episodic plume magmatism.

5.2 Introduction

Solid-state convection within Earth's mantle can be approximated as a fluid-dynamical process involving a number of complicating physical effects. These include internal heating from radioactive decay and bottom heating from the core, endothermic and exothermic phase changes, chemical heterogeneities, complex interior rheology and the complications arising from solidification and brittle failure at the surface boundary layer, i.e. tectonic plates. The interaction between these processes and features is difficult to investigate, particularly at the dynamical regime of Earth's mantle, which convects with an effective Rayleigh number (Ra) of $\approx 10^9$ (Davies 1999, Schubert et al. 2001). Numerical modeling is often the most feasible option for examining such problems. However, due to the complex and large-scale geometries involved and, despite the large computing clusters available to us today, thermal and chemical boundary layers have rarely been resolved at such vigor, in global convection models. By employing a modified and highly efficient extension to the three-dimensional (3-D) spherical mantle dynamics code TERRA (Baumgardner 1985, Bunge & Baumgardner 1995, Yang & Baumgardner 2000), our models achieve Rayleigh

numbers of $> 1 \times 10^9$. To the best of our knowledge, this is the first investigation of mantle convection phenomena at these vigors, in global 3-D spherical geometry.

The degree to which hot-spots remain fixed is a long-standing issue in global geophysics. Although some argue that hot-spots and mantle plumes are unrelated (e.g. Anderson 2000, Foulger et al. 2001), most accept that both are closely linked, with many believing that hot-spots represent the surface expression of deep rooted mantle plumes (e.g. Morgan 1972, Richards et al. 1989, Montelli et al. 2004). For the plume model to succeed, it needs to dynamically predict relatively stationary plumes for time-intervals of 50-100 million years (Myr) (Morgan 1983, Jurdy & Gordon 1984, Molnar & Stock 1987, Tarduno & Gee 1995). Both laboratory and numerical models have been presented that lead to plume fixity. Laboratory experiments generally require a layered system (Davaille et al. 2002, Jellinek & Manga 2002), while the majority of numerical models that yield steady plumes have been at low to moderate vigor (Bercovici et al. 1989, Zhong et al. 2000, Monnereau & Quere 2001). The highest vigor achieved, to date, in 3-D spherical geometry, was by Davies (2005), who showed that in models incorporating an increased lower mantle viscosity, it is possible to generate very steady, permanent plumes. This result was surprising. In two-dimensional (2-D) investigations, convection is shown to be steady at low Rayleigh numbers, but becomes more and more unsteady with increasing Rayleigh number (e.g. McKenzie et al. 1973, 1974, Jarvis 1984, Olson 1987, Christensen 1989). Steady convection has also been demonstrated at low vigor in 3-D spherical geometry (e.g. Bercovici et al. 1989, Ratcliff et al. 1996), therefore it was natural to expect 3-D flow to become unsteady at higher Rayleigh numbers. Further, since convective velocities in boundary layer theory of Rayleigh-Benard convection scale as $Ra^{\frac{2}{3}}$ (Turcotte & Oxburgh 1967), one might expect plume fixity to become weaker as the Rayleigh number increases. However, Davies (2005) showed that permanent plumes can arise in whole mantle convection models, at reasonably high vigors. With the debate

still raging as to whether or not hot-spots are the result of mantle plumes, it is of great importance to discover whether these conclusions hold at Earth-like vigor, which, for the purposes of this investigation, is taken as $Ra \approx 1 \times 10^9$. This was the motivation behind this study.

5.3 Methodology

5.3.1 Numerical Methodology

A modified version of the 3-D spherical mantle dynamics code, TERRA (Baumgardner 1985, Bunge & Baumgardner 1995, Yang & Baumgardner 2000), is employed to solve the incompressible, infinite Prandtl number conservation equations of mass (5.1), momentum (5.2) and energy (5.3), inside a spherical shell:

$$\nabla \cdot \mathbf{u} = 0 \quad (5.1)$$

$$\nabla \cdot \left[\mu \left(\nabla \mathbf{u} + (\nabla \mathbf{u})^T \right) \right] - \nabla p + \rho_0 g \hat{\mathbf{r}} \beta [T - T_0] = 0 \quad (5.2)$$

$$\frac{\partial T}{\partial t} + \nabla \cdot (T \mathbf{u}) = \kappa \nabla^2 T + \frac{H}{\rho c_p} \quad (5.3)$$

In the above equations, $\mathbf{u}$ is the fluid velocity vector, p denotes dynamic pressure, T absolute temperature, T_0 reference temperature (taken here as the surface temperature), ρ_0 reference density, t time, κ thermal diffusivity, g gravitational acceleration, β the coefficient of thermal expansion, c_p specific heat at constant pressure, H the heat generation rate per unit volume, μ dynamic viscosity (which can be a function of temperature and/or pressure) and $\hat{\mathbf{r}}$ the unit radial vector. These equations can be non-dimensionalized via the following relations:

$$x' = \frac{x}{d}; \quad t' = \frac{t\kappa}{d^2}; \quad T' = \frac{(T - T_0)}{\Delta T}; \quad u' = \frac{ud}{\kappa}; \quad Q' = \frac{Hd^2}{\kappa\Delta T\rho c_p}; \quad \mu' = \frac{\mu}{\mu_0}; \quad p' = \frac{d^2p}{\mu_0\kappa} \quad (5.4)$$

where ‘primes’ indicate the non-dimensional terms. Here, d is the thickness of Earth’s mantle, μ_0 is a reference viscosity (in this case, upper mantle viscosity) and ΔT is the

temperature drop from the core-mantle boundary (CMB) to the surface. Dropping the primes, Equations (5.1)–(5.3) become:

$$\nabla \cdot \mathbf{u} = 0 \quad (5.5)$$

$$\nabla \cdot \left[\mu \left(\nabla \mathbf{u} + (\nabla \mathbf{u})^T \right) \right] - \nabla p + RaT\hat{\mathbf{r}} = 0 \quad (5.6)$$

$$\frac{\partial T}{\partial t} + \nabla \cdot (T\mathbf{u}) = \nabla^2 T + Q \quad (5.7)$$

where Ra is the Rayleigh number, a non-dimensional parameter that quantifies the convective vigor of a system, being defined as:

$$Ra = \frac{\rho g \beta \Delta T d^3}{\kappa \mu_0} \quad (5.8)$$

It is important to note that in this study, the manner in which Rayleigh numbers are calculated depends upon the mode of heating. For a purely basally heated system, the Rayleigh number is calculated from Equation (5.8), being termed Ra_T . However, for internally heated cases, there is no fixed lower boundary temperature. As a result, the convective vigor is described by another term, Ra_H , which is given by:

$$Ra_H = \frac{\rho^3 g \beta H d^5}{k \kappa \mu_0} \quad (5.9)$$

where k is the thermal conductivity. Note that for the mixed-mode convection simulations included herein, which incorporate both basal and internal heating, values for both Ra_T and Ra_H are quoted, for completeness.

The adapted version of TERRA, introduced and validated in Chapter 4, allows one to solve these governing equations on a non-uniform grid. This leads to great flexibility and overcomes many of the code's limitations. A disadvantage to TERRA's original discretization scheme, derived by projecting the regular icosahedron onto the sphere, is the fact that element sizes and inter-nodal distances are dependent upon radius, with grid points denser at

the inner surface of the spherical shell when compared to the outer surface. To counteract this setup, the upper reaches of the grid are refined in the simulations presented here (upper 25% and 12.5% of the shell in low/moderate - $Ra_T \leq 1 \times 10^7$ - and high Rayleigh number - $Ra_T > 1 \times 10^7$ - simulations, respectively). This modification integrates well with the original scheme, rendering lateral resolution virtually identical at both boundaries. This is highly beneficial since both boundary layers can be simultaneously modeled at the same resolution, which was not possible with the original code. As a result, some of the simulations presented in this study achieve $Ra_T = 1 \times 10^8$ and $Ra_H = 1.4 \times 10^9$. They are therefore at superior vigor to the highest attained, to date, in global 3-D spherical geometry.

A further advantage to the modified code is its efficiency in terms of the CPU-time needed to solve a given problem. The solution algorithm is resourceful and highly practical, meaning that a rapid parameter space mapping can be performed. This allows us to investigate, systematically, the influence of Rayleigh number and viscosity structure upon the nature of upwelling structures, without expending vast amounts of CPU-time.

5.3.2 Model Configuration

We model isochemical, incompressible convection, employing the Boussinesq approximation, where the density is held constant ($4.5 \times 10^3 \text{ kg m}^{-3}$) other than in the buoyancy term of the momentum equation. For simplicity, we assume constant thermal expansivity ($2.5 \times 10^{-5} \text{ K}^{-1}$), thermal conductivity ($4.0 \text{ W m}^{-1} \text{ K}^{-1}$) and heat capacity ($1.13 \times 10^3 \text{ J kg}^{-1} \text{ K}^{-1}$). Mechanical boundary conditions are free-slip, while temperature boundary conditions are isothermal, set to 300 K at the surface and 3000 K at the CMB. In models that include internal heating, a volumetric heating rate of $4.0 \times 10^{-12} \text{ W kg}^{-1}$ is prescribed. Due to computational and numerical limitations, the simulations presented do not include a temperature dependent rheology. However, previous work suggests that its inclusion would

Parameter	Value	Units
Outer Shell Radius	6370	km
Inner Shell Radius	3480	km
Thermal Conductivity	4.0	$\text{W m}^{-1} \text{K}^{-1}$
Thermal Expansivity	2.5×10^{-5}	K^{-1}
Density	4.5×10^3	kg m^{-3}
Specific Heat at Constant Pressure	1.13×10^3	$\text{J kg}^{-1} \text{K}^{-1}$
Surface Temperature	300	K
CMB Temperature	3000	K
Acceleration Due to Gravity	10	ms^{-2}
Heat Generation	4.00×10^{-12}	W kg^{-1}
Clapeyron Slope (410 km)	0.0	MPa K^{-1}
Clapeyron Slope (660 km)	0.0	MPa K^{-1}

Table 5.1: Model parameters.

Case Number	Heating Mode	Ra_T	Ra_H	μ_1 (Pa s)	μ_2 (Pa s)	μ_2/μ_1
1	Basal	1×10^6	0	9×10^{22}	9×10^{22}	1
2	Basal	1×10^8	0	9×10^{20}	9×10^{20}	1
3	Mixed	1×10^8	1.4×10^9	9×10^{20}	9×10^{20}	1
4	Mixed	1×10^6	1.4×10^7	9×10^{22}	3.6×10^{24}	40
5	Mixed	1×10^7	1.4×10^8	9×10^{21}	3.6×10^{23}	40
6	Mixed	1×10^8	1.4×10^9	9×10^{20}	3.6×10^{22}	40
7	Mixed	1×10^7	1.4×10^8	9×10^{21}	4.5×10^{22}	5
8	Mixed	1×10^7	1.4×10^8	9×10^{21}	1.8×10^{23}	20
9	Mixed	1×10^7	1.4×10^8	9×10^{21}	4.5×10^{23}	50
10	Mixed	1×10^7	1.4×10^8	9×10^{21}	4.5×10^{24}	500

Table 5.2: The varying parameters of each case examined in this study. The table includes the heating mode and Rayleigh numbers, calculated based upon upper mantle viscosities (μ_1). The upper section of the table covers isoviscous simulations (i.e. $\mu_2/\mu_1 = 1$, where μ_2 is the lower mantle viscosity), while the second section includes models with a viscosity structure that consists of two distinct layers, with μ_2 having a relative viscosity of $40 \times \mu_1$. The lower section comprises simulations with variable μ_2/μ_1 .

not dramatically alter the basic planforms yielded (Zhong et al. 2000, Lowman et al. 2001).

For all cases, the initial condition consists of small-scale, random temperature perturbations.

We have verified that results are independent of this initial condition. Model parameters are summarized in Table 5.1.

The models presented fall into three distinct categories. Isoviscous simulations are examined first, at a range of Rayleigh numbers (cases 1-3). These are used for reference, allowing us to pinpoint the effects of more complex viscosity structures in later models. We then examine simulations that include a layered viscosity structure (cases 4-6), with the viscosity in the lower mantle 40 times greater than the upper mantle, as illustrated in Figure 5.1. Although




Figure 5.1: The variation in viscosity factor with depth, for the layered viscosity simulations examined herein.

the exact nature of the mantle's viscosity structure remains a subject of broad debate, the profile included here, to first order, is plausible. Indeed, such a lower-mantle viscosity increase has been suggested by several independent lines of evidence (e.g. Hager 1984, Ricard et al. 1992, Mitrovica & Peltier 1995, Ranalli 2001, Forte et al. 2002, Mitrovica & Forte 2004). The final set of simulations examined incorporate variable upper/lower mantle viscosity ratios (cases 7-10). Model parameters that differ between cases are listed in Table 5.2.

Cases at low/moderate Rayleigh numbers use approximately 30 million nodes. This yields a lateral spacing of around 28 km at the surface and CMB. Higher Rayleigh number runs use approximately 160 million nodes, which yields a lateral spacing of around 14 km at both boundaries. This increased resolution is necessary due to the finer features inherent to high Rayleigh number convection. In all results that follow, model time is normalized to 'Earth-time' by multiplying the ratio of model *RMS* surface velocity to that of Earth's plates (Gurnis & Davies 1986).

5.4 Results & Discussion

5.4.1 Isoviscous Convection

Snapshots of the thermal field from cases 1, 2 and 3 are shown in Figure 5.2. All have an isoviscous viscosity structure, although the Rayleigh number and heating mode differs between simulations (see Table 5.2). By examining Figures 5.2(a) and (b) it can be seen that the Rayleigh number strongly affects internal dynamics, with the scale of features becoming finer with increasing Rayleigh number and, hence, decreasing viscosity. This process can be easily understood. In classic Rayleigh-Benard systems, convection tends to homogenize the temperature field thereby avoiding a gravitationally unstable situation for the greater part of the layer. However, directly adjacent to the top and bottom boundaries the fluid temperature makes a rapid transition from the largely homogeneous interior value to the colder or hotter boundary value. These transitional regions are known as thermal boundary layers. In effect, convection confines gravitationally unstable parts of the fluid to these relatively thin boundary layers (e.g. Turcotte & Oxburgh 1967, Roberts 1979, Mitrovica & Jarvis 1987).

In a fluid stirred by vigorous convection, more material will be homogeneous, thus the boundary layers will become thinner. Although this leads to yet more gravitationally stable material it also causes sharper thermal gradients in the boundary layers and thus even larger heat fluxes in to and/or out of the layer. Therefore, the more vigorous the convective stirring, the greater the heat flow through the layer (Hewitt et al. 1980, Schubert & Anderson 1985, Bercovici et al. 2000).




Figure 5.2: The temperature distribution from a series of incompressible, isoviscous convection calculations with variable Rayleigh number and heating mode. Models are otherwise identical. (a) Case 1 - basally heated, $Ra_T = 1 \times 10^6$, (b) Case 2 - basally heated, $Ra_T = 1 \times 10^8$, (c) Case 3 - basal and internal heating, $Ra_T = 1 \times 10^8$, $Ra_H = 1.4 \times 10^9$. The upper images are global snapshots of the thermal field showing a radial surface just above the CMB and a cross-section, being surrounded by a latitude/longitude grid to give the reader a sense of the mantle's volume. The central and lower images represent radial slices ≈ 50 km below the lithosphere and at 2000 km depth respectively. The scale illustrates the temperature, away from the lateral average, i.e. the thermal anomaly. Note that in this figure, and those to follow, black/white regions displayed on radial surfaces represent sections of the model that are hotter/colder than the scale's upper/lower limit.

From here, fluid sinks (rises), thereby feeding vertical convective currents in the form of downwellings (upwellings), which are driven by the negative (positive) buoyancy of the cold (hot) boundary layer and are retarded largely by viscous stresses. With more vigorous convection characterized by sharper boundary layer thermal gradients, buoyancy effects are amplified and instabilities form with greater ease than is the case in less vigorous systems. Furthermore, the reduced viscosity of systems at higher vigor means that instabilities need not accumulate as great a mass before overcoming viscous stresses and sinking (rising). As a consequence, higher Rayleigh number convection is characterized by smaller, but more numerous instabilities (Turcotte & Oxburgh 1967). This accounts for the disparity observed between cases 1 and 2.

An important observation to make, relevant to both cases, is that in a purely basally heated, unit aspect ratio system, convection is naturally symmetric, with upwelling and downwelling currents retaining some basic symmetry. Indeed, apart from the sign of their thermal anomalies, downwellings and upwellings will generally be indistinguishable (Weinstein & Olson 1990, Bercovici et al. 2000). However, in spherical geometry, this is not strictly true, since upper and lower boundaries are of different dimensions; the upper surface has an area four times larger than the CMB. Consequently, although to first order these are symmetric systems, with boundary layers handling the same heat flux, the instabilities that develop are not truly symmetric.

If we now examine case 3, a simulation with both internal and basal heating (Figure 5.2c), we have a form of convection that is more complex. When compared to case 2, the planform differs significantly, displaying greater thermal heterogeneity. The reasons for this can be understood if one appreciates that the upper thermal boundary layer must handle the heat injected through the bottom as well as heat generated internally. To accommodate this,

the upper thermal boundary layer develops a larger temperature drop (to sustain a larger thermal gradient) than the lower thermal boundary layer. In this way, the top boundary layer has a larger thermal gradient than the bottom one, leading to more numerous cold downwellings when compared to hot upwellings. Thus, within this model, we are seeing a top thermal boundary layer with a large temperature drop across it feeding downwellings which dominate the overall convective planform (e.g. Davies 1999, Bercovici et al. 2000, Schubert et al. 2001). This simple picture of internally heated convection is in keeping with the idea that large scale mantle circulation is driven by downwellings fed by an intense thermal boundary layer, while active upwellings are relatively weak and/or few in number (e.g. Davies & Richards 1992).

5.4.2 Layered Viscosity Convection

The thermal profiles yielded by cases 4, 5 and 6 are displayed in Figure 5.3. Excluding the Rayleigh number, each simulation is identical, incorporating both basal and internal heating, in addition to a layered viscosity structure. The increase in lower mantle viscosity completely transforms the planform, with greater thermal heterogeneity displayed than in isoviscous cases (note the differing scales of Figures 5.2 and 5.3). Long, linear sheet-like downwellings dominate upper mantle convection, as opposed to the isolated downwelling plumes previously observed. These widen substantially with depth. Indeed, as they enter the high viscosity lower mantle, they break up into large diffuse blobs, which is consistent with the findings of previous investigations (Hansen et al. 1993, Zhang & Yuen 1995, Bunge et al. 1996, 1997, Monnereau & Quere 2001, Lowman et al. 2004).




Figure 5.3: As in Figure 5.2, but for a series of simulations incorporating both basal and internal heating, in addition to a layered viscosity structure. (a) Case 4 - $Ra_T = 1 \times 10^6$, $Ra_H = 1.4 \times 10^7$ (b) Case 5 - $Ra_T = 1 \times 10^7$, $Ra_H = 1.4 \times 10^8$ (c) Case 6 - $Ra_T = 1 \times 10^8$, $Ra_H = 1.4 \times 10^9$. Note the different scales employed here when compared to Figure 5.2.




Figure 5.4: A series of snapshots from simulations with variable upper to lower mantle viscosity ratios. (a) Case 7 - $\mu_2/\mu_1 = 5$; (b) Case 8 - $\mu_2/\mu_1 = 20$; (c) Case 9 - $\mu_2/\mu_1 = 50$; (d) Case 10 - $\mu_2/\mu_1 = 500$. Each image shows a radial section just above the CMB and an iso-surface, representing regions of the model that are 500 K hotter than average for their depth. Again, the scale illustrates the thermal anomaly. Simulations with increased μ_2/μ_1 are characterized by fewer, but larger and more robust upwelling plumes.

Although this is true for all layered cases, the morphology of downwellings changes with Rayleigh number. For case 4, downwellings are formidable, being strong and coherent and displaying large thermal anomalies. However, with increasing vigor (cases 5 and 6) downwellings become thinner and weaker, but more numerous. Their thermal anomalies decrease significantly, which is consistent with the behavior exhibited in isoviscous simulations and fits in perfectly with simple boundary layer theory (e.g. Turcotte & Oxburgh 1967).

In all cases, we find that the downwellings ultimately reach the CMB. As they descend, they push hot material down and away. This process forces hot mantle to collect in a cellular pattern at the CMB. Over time, the continued supply of downwelling material squeezes hotter fluid into a network of narrow linear ridges. The return upward flow is focused at the junctions of these ridges, as cylindrical upwellings/plumes. As with downwelling structures, the morphology of upwellings changes dramatically with Rayleigh number; higher vigor systems are characterized by narrower and weaker plumes, although it is interesting to note that their number remains reasonably constant. A similar behavioral pattern is observed when decreasing the ratio between lower and upper mantle viscosities, μ_2/μ_1 (see Figure 5.4). Simulations with decreased μ_2/μ_1 are characterized by smaller and weaker plumes, although this is to be expected; by decreasing μ_2/μ_1 , one increases the effective Rayleigh number. However, when decreasing μ_2/μ_1 , the number of plumes increases. Such a pattern is not observed when increasing the Rayleigh number, suggesting that the ratio between lower and upper mantle viscosities is the predominant control upon the number of upwellings in these systems.

Perhaps the most striking feature of these models is that plumes are extremely coherent and, as illustrated by the temporal evolution plots of Figures 5.5 (case 4) and 5.6 (case 5), for the lower vigor cases, appear almost stationary and anchored in their positions. Indeed,




Figure 5.5: A series of thermal profiles, displaying the temporal evolution of case 4 ($Ra_T = 1 \times 10^6$, $Ra_H = 1.4 \times 10^7$). Snapshots are spaced ≈ 50 Myr apart. Each one shows a radial surface just above the CMB and a hot iso-surface, representing regions of the model that are 500 K hotter than average for their depth. They are surrounded by a latitude/longitude grid to give the reader a sense of the mantle's volume. The most prominent features are hot upwelling plumes, which are stationary, robust and coherent.




Figure 5.6: As in Figure 5.5, but for case 5 ($Ra_T = 1 \times 10^7$, $Ra_H = 1.4 \times 10^8$). Again, the most prominent features are hot upwelling plumes, which are strong and powerful.

once formed, they typically survive the remaining length of the simulation, which exceeds 500 Myr in both cases. Over this time, subtle drifting is observed in the upwellings of case 5 and, on occasions, smaller plumes coalesce. Nonetheless, the larger and dominant plumes remain remarkably stable. With case 6 however, the planform is more complex. Although powerful steady plumes remain a feature of this simulation, they do not dominate internal dynamics as in previous cases (see Figure 5.7). Indeed, a large number of mobile and ephemeral plumes are observed, which range in intensity and drift velocities. New plumes form throughout the simulation and old ones are seen to fade, many of which eventually cease. In addition, nearby plumes are seen merging together over time. Those plumes that do remain steady appear to alter in intensity as the simulation evolves. For example, the plume marked with a red cross in Figure 5.7(a) decreases in intensity for 80 Myr, from (a) to (c). Over the next 120 Myr however, it appears to increase in strength, with a large plume head visible in Figure 5.7(f), which is approximately twice the size of the head observed in Figure 5.7(a) and five times the size of that observed in Figure 5.7(c). It is not clear whether such behavior has been observed in previous purely thermal models. This therefore is an important observation that will be discussed later. However, the primary result of these simulations is that upwelling behavior varies with Rayleigh number; plume fixity decreases with increasing vigor, i.e. in moving from case 4 through to case 6. Although steady plumes are observed in case 6, they coexist with mobile and ephemeral upwellings and, hence, the conclusions of previous studies (e.g. Davies 2005), with large, steady plumes dominating mantle dynamics, do not hold at truly Earth-like vigor.

The obvious and most simple explanation for this shift in behavior is that the upwellings generated in case 6 are part of a more vigorously convecting system. As a consequence, they are swept around and occasionally broken up by large-scale background flow, resulting in shorter-lived plumes with decreased fixity. However, this straightforward explanation

does not tell the whole story. As stated previously, while descending through the mantle downwellings force hot material down and away, resulting in numerous relatively small, quasi-circular patches of cold material at the CMB. As these are supplied with more and more downwelling material, they push hotter fluid sideways, such that it collects at ridges, which form when hot material is brought together from opposite directions. Plumes occur at locations where these ridges meet. This depiction of the system's dynamics demonstrates how downwellings control and manipulate the upwellings. With downwellings becoming smaller and less powerful with increasing vigor, their ability to control and induce movement in hot material decreases, which could account for the decreased upwelling fixity in higher vigor systems. However, since upwellings are also weaker in such systems, downwelling potency cannot solely be the cause of the disparity observed between cases. The complete explanation is more subtle.

While it is true that downwellings dictate the locations of upwellings, the upwellings themselves also control the positioning of downwellings. This is particularly apparent in lower vigor simulations where plumes are strong, coherent and extremely powerful. Regions where they impact at the surface become too hot and hence, too buoyant, to sink. This self-organization makes the planform exceptionally stable and also allows the upwellings to become more robust, building up greater resistance to the background flow with time, thereby becoming anchored in their positions. However, with increasing vigor, upwellings become noticeably thinner and weaker, as is clear from Figure 5.3(c), where plumes are very faint in the upper mantle and at the surface. As a consequence, the upwellings are unable to control the location of downwellings to the same extent; the self-organization observed in lower vigor cases is not as apparent. Accordingly, plumes cannot build up such a resistance to the background flow and, hence, their fixity decreases.




Figure 5.7: As in Figure 5.5, but for case 6 ($Ra_T = 1 \times 10^8$, $Ra_H = 1.4 \times 10^9$), with snapshots spaced ≈ 40 Myr apart. We see that even at Earth-like vigor, steady plumes remain a feature of convection. However, they are accompanied by a large number of mobile and ephemeral plumes, which range in intensity and drift velocities. New plumes form throughout the simulation and old ones are seen to fade, many of which eventually cease. In addition, smaller plumes are seen merging together over time.

This line of reasoning also accounts for the altering intensity displayed by the stable plumes of case 6. We know that downwellings control the location of upwellings, but plume strength and vigor is also dictated by the supply of cold downwelling material to the CMB. With a consistent supply of material to this region, upwellings will display what is essentially a regular intensity through time, as indeed occurs in case 4 and, fundamentally, case 5. However, in less-organized systems, such as case 6, which are characterized by an irregular downwelling pattern, upwellings will show varying intensities at different stages of the simulation. At times of high supply of cold material to the CMB, hot fluid will be pushed into ridges more rapidly and, hence, upwellings will be stronger and more vigorous. At times of low supply, upwellings are likely to be weaker.

5.5 Discussion

This discussion will focus upon case 6, a layered viscosity simulation at a Rayleigh number of greater than 1×10^9 , examining its implications for convection within Earth's mantle. The upwelling characteristics displayed in this model are somewhat similar to the patterns observed on Earth. As noted previously, numerous analyses of hot-spot track and plate motion directions have established that hot-spots remain relatively stationary for time intervals of 50-100 Myr (e.g. Morgan 1983, Jurdy & Gordon 1984). Steady plumes that could account for such surface observations arise in our model, with their stationarity shown to strongly depend upon an increased lower mantle viscosity. Such reduced mobility has been observed in previous studies (e.g. Zhang & Yuen 1995, Bunge et al. 1996, Lowman et al. 2004) and is easy to understand. The increased viscosity with depth reduces the effective Rayleigh number of the lower mantle and, hence, the characteristic time scales for flow become considerably longer at depth. As a consequence, the time-scales for instability are

longer in the lower mantle. Accordingly, dynamical structures originating at depth, such as the upwelling plumes observed herein, are extremely persistent.

The stationary plumes that arise in this model probably have a longer lifespan than Earth's hot-spots. One of the oldest observable hot-spots on Earth is Hawaii. The Hawaiian-Emperor track, a 4000 km long chain of volcanic islands and seamounts, running from near the Aleutian Islands to the active Kilauea volcano, represents at least 110 Myr of volcanism (Tarduno & Gee 1995, Schubert et al. 2001). Although our plumes are not precisely fixed, which is consistent with the findings of Steinberger & O'Connell (1998), some remain 'stationary' for periods that exceed 200 Myr. However, the geological record suggests that such a situation does not arise on Earth. Although several factors could potentially contribute to this disparity, we argue that the major control is the influence of tectonic plates at the surface, which break the symmetry and self-organization of upwelling and downwelling currents, thereby reducing the probability of plumes remaining stationary for extended periods of time.

In addition to the steady plumes that arise, mobile and ephemeral plumes are a feature of our model. Indeed, considered as a whole, the upwellings display many morphological characteristics that are consistent with observations for Earth's hot-spots. These include a range of longevities, drift velocities, intensities and sizes, in addition to an uneven distribution across the globe. Such features are not observed in the less vigorous models of Davies (2005), nor in the lower vigor simulations examined herein, suggesting that this behavior is a function of the Rayleigh number. At higher vigors, upwellings can be broken up and swept-around by large-scale background flow, resulting in shorter-lived plumes with decreased fixity.

Although most hot-spots, such as Hawaii and Reunion, display a remarkably uniform age

progression, others are variable; some display inconsistencies in the progression of their ages, while others have simultaneously active volcanism along the track. Some have even been known to produce either segmented tracks or, on occasions, a short pulse of activity. Even the Hawaiian hot-spot has exhibited large variations in strength as evidenced by variations in the rates of melt production and swell formation (Davies 1992). If hot-spots are indeed related to mantle plumes, this would suggest that a wide-range of upwelling structures occur within Earth's mantle, as indeed has been suggested by Courtillot et al. (2003) and Davies & Bunge (2006). The model presented here, with variable plume like features coexisting, is therefore well-suited to explaining observations on Earth.

Perhaps the most dramatic example of variations in hot-spot intensity is found in the record of flood basalts, which often show evidence for distinct episodes of eruptive activity, at intervals ranging from a few million to tens of millions of years. A number of previous models have been proposed to explain such episodic plume magmatism. These include thermal interactions and plume head separation at the 660-km discontinuity (Bercovici & Mahoney 1994, Schubert et al. 1995) and a secondary instability in non-newtonian fluids (van Keken 1997). Such phenomena could explain the second major magmatic event in a sequence, since a secondary plume head is formed at the original plume conduit. However, they cannot account for multiple events with variable strengths at irregular time intervals. Other possible causes of such 'pulsing' behavior include solitary waves traveling along the low viscosity, thermal plume conduit (Olson 1987) and the interaction of magma transport and lithospheric flexure (Hieronymus & Bercovici 1999). However, it is unlikely that these mechanisms could produce the large volumes of melt recorded in discrete flood basalt events. Multiple plume sources or a single dismembered plume have been considered as the cause of the high flux and diverse geographic distribution of magmatism generated by the Kerguelen plume 120–95 Myr ago (Coffin et al. 2002), but these are unlikely to be a global

feature. Lin & van Keken (2005) have argued that the processes leading to large igneous provinces are more complicated than the purely thermal, single-stage plume model suggests. They convincingly demonstrate that the multiple episodes of major flood basalt eruptions could result from the interaction between thermal and compositional buoyancy forces, i.e. thermo-chemical convection. The characteristic timescales of secondary instabilities and the variation in plume intensity in their models are compatible with observations on Earth.

Although composition doubtless plays a role in shaping plume strength and intensity on Earth, our simulation suggests that downwellings, in the form of subducting slabs, might also be important. In our model, upwelling intensity is dictated by the supply of cold downwelling material to the CMB. As they descend, downwellings push hot material down and away, forcing hot mantle to collect in a cellular pattern at the CMB. Over time, the continued supply of downwelling material squeezes hotter fluid into a network of narrow linear ridges. The return upward flow occurs where these ridges meet, as cylindrical upwellings/plumes. At times of high (low) supply of downwelling material to the CMB, hot material is swept into these ridges rapidly (slowly), contributing to an increased (decreased) plume buoyancy. Therefore, rather counter-intuitively, plume intensity is strongest (weakest) during, or shortly after, periods of heavy (light) pooling of cold material at the CMB. Within Earth's mantle, it is highly unlikely that steady-state and uniform plume source conditions will arise. Here, plumes coexist with larger-scale flows associated with, for example, plate motions. Since subduction is a time-dependent, asymmetric process (e.g. Silver & Behn 2008, Goes et al. 2008), were our conclusions correct, they would suggest that subducting slabs might be partially responsible for episodic plume magmatism (Parman 2007, Pearson et al. 2007).

The dominance of downwellings and their control upon internal dynamics also suggests that

slabs might account for other plume characteristics, most notably, their positioning. On Earth, subducting slabs could be moving broad hot-zones along the CMB, leading to sharp variations in lateral seismic velocity (Ni et al. 2002, Thomas et al. 2002) and even forming super-swells (Thompson & Tackley 1998, Davaille 1999). This picture of internal dynamics would imply that upwelling features would tend to only be found away from regions of the CMB that have recently suffered impacting of subducting slabs, as suggested by Tan et al. (2002) and Davies (2005). Indeed, it is argued that such behavior is observed (Anderson 1998, Chase 1985, Richards & Engebretson 1992). The majority of hot-spots, believed by many to be the surface manifestations of upwelling plumes, occur in regions of geoid high (Chase 1979, Stefanick & Jurdy 1984) and negative seismic velocity anomalies. The geoid high and seismic anomalies have been argued to arise from regions of the mantle without recent subducted material.

While this investigation has examined several aspects of high-vigor convection there is room for further work. The models presented are realistic in many ways. However, several important aspects of Earth's mantle have not been addressed, the most prominent of which are a temperature dependent rheology, chemical heterogeneities and tectonic plates at the surface. With regards to a temperature dependent rheology, due to potential non-linear feedbacks it is difficult to predict exactly how its inclusion would alter our results (e.g. Lowman et al. 2001, Korenaga 2005). Furthermore, it is very challenging to model large lateral viscosity variations robustly, with the best simulations to date in 3-D spherical geometry attaining Rayleigh numbers a few orders of magnitude below the dynamical regime of Earth's mantle (Zhong et al. 2000). Although such rheological properties will doubtless have an affect upon the dynamics of convection, as indeed will the inclusion of chemical heterogeneities (e.g. McNamara & Zhong 2005, Nakagawa & Tackley 2006), many characteristics of the simulations presented here remain representative of processes within Earth's

mantle. Nonetheless, modeling and constraining the temperature dependence of mantle viscosity as well as developing efficient ways to accurately track chemical heterogeneities, should be one focus of future work. In addition, since the hot-spot characteristics observed on Earth are hypothesized to occur due to the combination of deep plumes and plate dynamics (e.g. Tackley 2006), we need to better address the surface in our models and also, include continents, if we are to truly investigate such phenomena. This point, regarding the importance of the surface, is emphasized in the following chapter.

5.6 Conclusions

Strong, stationary and coherent plumes have been shown to develop at low to moderate vigors and in layered systems. Such configurations have also been observed in higher-vigor simulations incorporating a layered viscosity structure. However, due to computational limitations, prior to this study, such phenomena had not been examined at Rayleigh numbers of greater than $\approx 1 \times 10^8$. Earth's mantle is believed to convect at a Rayleigh number of order 10^9 .

Owing to a highly-efficient, multi-resolution version of the 3-D spherical mantle dynamics code TERRA, the simulations presented herein replicate such Rayleigh numbers. They are therefore at superior vigor to the highest achieved, to date, in global 3-D spherical geometry. Results demonstrate that at these vigors, steady plumes do arise. However, they do not dominate the planform as in lower vigor cases; they coexist with mobile and ephemeral plumes, which display wide-ranging characteristics. Since hot-spots, believed by many to be the surface expression of deep mantle plumes, are known to display a range of strengths, sizes and drift velocities, the variable upwelling pattern observed in our model is somewhat representative of observations on Earth. Those plumes that do remain steady

vary in intensity throughout the simulation, strengthening and weakening over time. We argue that such behavior is caused by an irregular supply of cold material to the core-mantle boundary region, which could point towards subducting slabs being partially responsible for episodic plume magmatism on Earth.

As is verified by the lack of publications at these high vigors, calculations of this nature would be unfeasible were it not for the modified code. Although our conclusions are significant, perhaps the most important lesson to take from this investigation arises from the methodology. As a community, we must develop and implement advanced, multi-resolution numerical schemes if the recent progress made in geodynamical modeling is to continue.

Chapter 6

The Importance of Earth's Surface Boundary Condition to Mantle Convection Simulations

6.1 Abstract

A modified and benchmarked version of the 3-D spherical mantle dynamics code TERRA, which allows for variable resolution radially, is employed to investigate the importance of Earth's surface boundary condition in mantle convection simulations. Three different mechanical boundary conditions are modeled (free-slip, rigid and kinematic with imposed plate motion history), in otherwise identical systems. Results demonstrate that the surface has a profound control upon internal dynamics, manifesting itself not only in convective structures, but also in thermal profiles, Nusselt numbers and velocity patterns. Since the majority of geodynamical simulations incorporate a surface condition that is not at all representative of Earth, this is a worrying, yet important conclusion. By failing to address

the surface appropriately, geodynamical models, regardless of their sophistication, cannot be truly applicable to Earth.

6.2 Introduction

Plate tectonics, the theory that underpins most of modern geology, is widely-regarded as the prime surface expression of mantle convection. Yet, surface deformation with piecewise uniform motion that would approximate tectonic plates is not easily obtained in mantle convection simulations. The general approach to incorporating plates within mantle convection models has been to specify some component of the system, such as surface velocities, weak zones or faults, while solving for the dynamics of the remainder of the system (Gurnis & Hager 1988, Davies 1989, Gable et al. 1991, Weinstein & Olson 1992, King & Hager 1994, Zhong & Gurnis 1995, Bunge et al. 1998, Lowman et al. 2001). Such methods have their strengths and weaknesses and have yielded important insights into various convection phenomena. However, as of yet, no single model can spontaneously reproduce true plate-tectonic behavior.

The difficulty in accounting for plates in numerical simulations may be explained if plates are a self-organized system that coordinates mantle convection, rather than vice versa. Upper mantle convection patterns should then be regarded as the result, not the cause, of plate tectonics. There are several proponents to such 'top-down' mechanisms (Anderson 2001) which either disregard convection, or focus upon pre-existing plates moved by forces such as slab-pull and ridge push (Forsyth & Uyeda 1975).

Whether changes in plate boundary forces cause shifts in plate motions or whether both changes in plate boundaries and plate motions are a response of the plate system to mantle convection is an important question. Records from plume related hot-spot volcanism suggest

that global plate motions remain relatively constant for extended periods of time. However, these periods are separated by episodes of abrupt transition, the most prominent example of which is recorded at the great bend in the Hawaiian-Emperor seamount chain (Molnar & Stock 1987, Norton 1995, Lithgow-Bertelloni & Richards 1998, Tarduno et al. 2003, Steinberger et al. 2004, Finn et al. 2005, Wessel et al. 2006, Whittaker et al. 2007). This fundamental and peculiar feature of plate tectonics is very difficult to understand. Plate motions derived purely from internal buoyancy forces can only evolve over many millions of years i.e. the timescales over which such buoyancy forces develop in the mantle. Thus, it is believed that sudden changes in plate motions cannot be explained by such models, which is a major line of evidence used to advocate the aforementioned 'top-down' mechanisms. Recent work, however, has suggested that hot mantle accumulating beneath plates can be swept toward and annihilate downwellings, leading to a relatively rapid adjustment in plate driving forces (Lowman et al. 2001, King et al. 2002). In either case, as noted by Bercovici (2003), the *a priori* existence of the plates precludes any attempt to explain how they are generated and evolve in the first place. The answer to this fundamental question of how plates arise from, or perhaps, cause mantle convection, has been the focus of an active, yet technically difficult field of research. Nonetheless, significant progress has been made and the general consensus now emerging is that plates and convection are one coupled system, which closely interact (Tackley 2000*a,b*, Lowman et al. 2001, Bercovici 2003).

Over the same period, major advances in our understanding of the complex Earth system have been made, with computational geodynamics connecting aspects of mantle convection theory with geophysical, geochemical and geological observations. A threshold has been crossed in high-performance computing, allowing 3-D spherical mantle convection simulations to achieve the resolution necessary to capture 'Earth-like' Rayleigh numbers, as demonstrated in Chapter 5. Rapid advances have also been made in modeling complex

rheologies. These developments hold great potential for interdisciplinary research in plate tectonics and mantle convection. However, the fact that convection models fail to reproduce plate-like behavior at the upper boundary is proving to be a major obstacle to such collaboration.

Several studies have demonstrated that, in mantle convection simulations, the upper boundary condition has a strong influence on the dynamics of the system. For example, the work of Tan et al. (2002) and Davies (2005), in both regional and full spherical geometry, as well as the simulations presented in Chapter 5, suggests that subducting slabs modulate plume formation, with plumes tending to develop at the edge of slabs. In areas of the core mantle boundary (CMB) free of slabs, plume formation is frequent while the basal thermal boundary layer is thin. However, in areas beneath slabs, the basal thermal boundary layer is thicker and plume formation infrequent. Indeed, it is argued that such behavior is observed (Anderson 1998, Chase 1985, Richards & Engebretson 1992). The majority of hot-spots, believed by many to be the surface manifestations of upwelling plumes, occur in regions of geoid high (Chase 1979, Stefanick & Jurdy 1984) and negative seismic velocity anomalies. The geoid high and seismic anomalies have been argued to arise from regions of the mantle without recent subducted material.

Lowman et al. (2004) further examine the influence of plate-like surface motion on upwelling dynamics through 3-D Cartesian models, concluding that plume positioning and formation is affected by the influence of plate-scale flow. One would expect such effects to be amplified in spherical geometry, where the variation in radius between inner and outer surfaces would become a key factor. Clearly, these studies, amongst others (Bunge et al. 1998, Conrad & Hager 1999, King et al. 2002, Davies & Bunge 2006, Quere & Forte 2006, Nettelfield & Lowman 2007), suggest that the surface exerts a major control on the evolution of man-

the convection, yet the mechanical upper boundary conditions utilized in the majority of geodynamical simulations are not at all representative of Earth.

The motivation behind this chapter therefore is to emphasize the importance of the surface boundary condition to mantle convection simulations. Even though previous studies have touched upon this topic, as of yet, no single study has looked at the problem in a general sense. Three distinct mechanical boundary conditions are simulated, i.e. stress-free/free-slip, rigid and kinematic, in otherwise identical 3-D spherical models. Results, as expected, illustrate that the surface plays a central role in shaping the dynamics of the interior, highlighting the fact that we, as a geodynamical community, must begin to better address the upper boundary condition. As stated by Tackley (2000*c*), it is of fundamental importance to improve our understanding of plate boundary processes so that the relevant material properties can be included in numerical models. This would allow plates to arise self-consistently from the constitutive equations, rather than being imposed by the modeler. Such capabilities would, without doubt, yield unprecedented opportunities for true interdisciplinary research.

6.3 Methodology

6.3.1 Numerical Methodology

A modified version of the 3-D spherical mantle dynamics code, TERRA (Baumgardner 1985, Bunge & Baumgardner 1995, Yang & Baumgardner 2000), is employed to solve the infinite Prandtl number conservation equations of mass, momentum and energy, inside a spherical shell. The adapted version, introduced and validated in Chapter 4, allows one to refine and solve on a radially non-uniform grid. Consequently, complex problems can be solved efficiently, both in terms of RAM and CPU-time.

For the purposes of this investigation, the domain is discretized into a fine upper mantle (the outer 25% of the shell) with coarser elements elsewhere, ensuring sufficient resolution to capture the upper thermal boundary layer. Such a setup counteracts the original discretization inherent to TERRA, which is characterized by superior resolution in the lower mantle. In addition, since the viscosity of Earth's mantle increases significantly with depth (Hager 1984, Ricard et al. 1992, Mitrovica 1996), one would expect upper mantle convection to be dominated by fine-scale activity, with longer wavelength features more prevalent in the lower mantle. The advocated scheme therefore fits in perfectly with such behavior, yielding great benefits for the investigations presented herein. Furthermore, the solution algorithm allows one to perform a rapid parameter space mapping. In essence, although the upper thermal boundary layer is simulated at high resolution, the overall CPU-time is not significantly greater than that for a run at lower resolution.

6.3.2 Model Configuration

The influence of three different surface mechanical boundary conditions is investigated, namely free-slip (case 1), rigid (case 2) and kinematic (case 3). Models are otherwise identical, with the mechanical boundary condition at the CMB always free-slip. For case 3, the history of plate motion is prescribed as a time-dependent velocity boundary condition, starting in the mid-Mesozoic, 119 million years ago (Ma) (Lithgow-Bertelloni & Richards 1998). By forcing the surface with plate velocities, one essentially anticipates that the dynamics of the remainder of the system evolve to be somewhat consistent with Earth. Such an approach therefore provides an avenue towards a quasi Earth-like boundary condition in this study. It should be noted that although the criticism made of the kinematic approach is valid (i.e. velocity boundary conditions do substantial work on the system), it does not incorrectly alter the evolution of a calculation when compared to a model having plates

generated self-consistently (Han & Gurnis 1999). Such a scheme is therefore ideal for the present study.

The runs presented use approximately 30 million nodes, which, bearing in mind the modified discretization, yields a lateral spacing of around 28 km at the surface and CMB. The radial spacing, finer in the upper 25% of the shell, is also reduced close to the CMB to accurately capture the lower thermal boundary layer. The scale of features is controlled by the viscosity, with upper mantle viscosity set at 1×10^{22} Pa s. This choice allows us to be confident of both resolved calculations and a reasonably high-convective vigor (whole mantle $Ra = 2.5 \times 10^6$, calculated based upon basal heating). Our calculations easily surpass the rule of thumb of 5 points per high gradient region quoted by Lowman et al. (2004). We simulate compressible convection, employing the Murnaghan equation of state (Murnaghan 1951, Bunge et al. 1997). This leads to a coefficient of thermal expansion that fits experimental observations (Chopelas & Boehler 1992), decreasing with depth from $\approx 4 \times 10^{-5} K^{-1}$ at the surface to $\approx 1.25 \times 10^{-5} K^{-1}$ at the CMB, as illustrated in Figure 6.1(a). Reference density (Figure 6.1b) increases monotonically from $3.5 \times 10^3 \text{ kg m}^{-3}$ at the surface to $5.5 \times 10^3 \text{ kg m}^{-3}$ at the CMB, in good agreement with the preliminary reference Earth model - PREM (Dziewonski & Anderson 1981). Heat capacity and thermal conductivity are constant, with the former being set to $1.13 \times 10^3 \text{ J kg}^{-1} \text{ K}^{-1}$ (Stacey 1992) and the latter to $4.0 \text{ W m}^{-1} \text{ K}^{-1}$. An internal heating rate of 4.0×10^{-12} is imposed, which is just below the chondritic value. Other model parameters are listed in Table 6.1.

All simulations incorporate depth-dependent viscosity, with a structure similar to that proposed by Mitrovica & Forte (2004), derived via joint inversion of convection and glacial isostatic data. The convection related observables included satellite-derived free-air gravity harmonics, the geodetically inferred excess ellipticity of the CMB and tectonic plate




Figure 6.1: Variation in (a) the coefficient of thermal expansion, (b) density and (c) viscosity, as a function of depth through the mantle.

Parameter	Value	Units
Thermal Conductivity	4.0	$\text{W m}^{-1} \text{K}^{-1}$
Specific Heat at Constant Pressure	1.13×10^3	$\text{J kg}^{-1} \text{K}^{-1}$
Surface Temperature	300	K
CMB Temperature	3000	K
Acceleration Due to Gravity	10	ms^{-2}
Heat Generation	4.00×10^{-12}	W kg^{-1}
Clapeyron Slope (410 km)	0.0	MPa K^{-1}
Clapeyron Slope (660 km)	0.0	MPa K^{-1}

Table 6.1: Model parameters.

motions. It is important to note that when compared to the profile of Mitrovia & Forte (2004), our variations over the mantle are reduced to ensure computational stability. As illustrated by Figure 6.1(c), the simulated viscosity profile shows a significant ($30\times$) increase from the upper mantle to a high-viscosity peak at 2000 km depth, followed by a reduction toward the core-mantle boundary. A low-viscosity region at 660 km is not incorporated. The lithosphere is modeled by specifying a highly viscous layer, of thickness 150 km, at the top of the mantle to emulate the high viscosity of the mantle's cold upper thermal boundary layer. Due to computational limitations, models do not include a temperature dependent rheology.




Figure 6.2: Plate boundary maps for Mesozoic and Cenozoic plate stages. (a) 119-100 Ma; (b) 94-84 Ma; (c) 74-64 Ma; (d) 48-43 Ma; (e) 43-25 Ma and (f) present day (10-0 Ma). Arrows indicate the direction of plate motion. Their length is proportional to the plate speed. AF denotes Africa; AN, Antarctica; AR, Arabia; AU, Australia; CA, Caribbean; CO, Cocos; CR, Chatham Rise; EU, Eurasia; FA, Farallon; IN, India; IZA, Izanagi; KU, Kula; NA, North America; NZ, Nazca; PA, Pacific; PL, Philippines; PH, Phoenix and SA, South America. In the Mesozoic, the ancient Izanagi (IZA), Farallon (FA), Kula (KU) and Phoenix (PH) plates occupy most of the Pacific Basin. Today these plates have largely disappeared. Reproduced from Bunge et al. (2002), which was compiled from Lithgow-Bertelloni & Richards (1998).




Figure 6.3: Comparison between models of mantle convection with different surface conditions. (a) Case 1, free-slip condition; (b) Case 2, rigid condition; (c) Case 3, imposed plate velocity kinematic condition. The left and central images represent the thermal anomalies at 150km and 2200km depth respectively. For case 3, plate boundaries are outlined. The scales illustrate the temperature, away from the lateral average, i.e. the thermal anomaly. Note that black regions displayed in case 1 represent sections of the model that are hotter than the scale's upper limit. Images on the right are snapshots of the thermal field, across the whole sphere. Each one shows a radial surface just above the CMB and a hot iso-surface, representing regions of the model that are 400K hotter than average for their depth. Images are surrounded by a latitude/longitude grid to give the reader a sense of the mantle's volume.

For cases 1 and 2 the initial condition is made up of small-scale, random temperature perturbations. For case 3, an initial condition is generated by driving the surface with the plate motions of the Mid-Cretaceous (see Figure 6.2a) until a quasi steady-state of convection is achieved. We account for the reduced convective vigor of our models by scaling the root-mean-square (*RMS*) plate velocities so that on average, imposed plates neither speed up nor slow down the flow from what would normally be generated with a free-slip top-surface (Bunge et al. 2002). It is important to note that due to the reduced convective vigor of our models when compared to Earth, absolute times are not useful. Instead, we express time in terms of transit times (Gurnis & Davies 1986), where one transit time is defined as the ratio of mantle depth to model *RMS* surface velocity. Thus, one transit time for Earth, assuming an *RMS* plate velocity of 4.4cm yr^{-1} , corresponds to about 65Ma. In all calculations that follow, model time is normalized to 'Earth-time' by multiplying by the ratio of Earth transit time to model transit time, or equivalently, by multiplying by the ratio of model *RMS* surface velocity to that of Earth's plates.

6.4 Results & Discussion

Figure 6.3 displays snapshots of the thermal field from all simulations. Excluding the upper boundary condition these models are identical, yet, the convective planforms display significant variations. This illustrates, beyond any doubt, that the surface condition has a profound influence upon internal dynamics.

Beginning with results from the free-slip simulation, i.e. case 1, we see that the planform is dominated by the sinking of the cold upper boundary layer as sheet-like downwellings. These widen substantially with depth and break up into large diffuse blobs as they enter and traverse the high viscosity lower mantle. These patterns are consistent with previous

investigations (Hansen et al. 1993, Zhang & Yuen 1995, Bunge et al. 1996, Monnereau & Quere 2001, Lowman et al. 2004), providing further evidence to support the conjuncture that increased viscosity broadens the cold thermal anomalies associated with sinking sheets. This results in a more diffuse pattern of heterogeneity than is the case with uniform properties.

As downwellings descend through the mantle, hot material is forced down and away. This behavior results in numerous relatively small, quasi-circular patches of cold material at the CMB. As these are supplied with more and more material, they squeeze the hotter fluid into a network of narrow linear ridges. The return upward flow is focused at the junctions of these ridges, as strong cylindrical upwellings/plumes. The tendency is for nearby plumes to join together over time and, hence, they decrease in number as the simulation evolves, with the final planform exhibiting only four plumes. These however, are extremely coherent, almost stationary and appear anchored in their position. Indeed, once formed, they typically survive the remaining length of the simulation, as illustrated by Figure 6.4.

While it is true that in these simulations downwellings control and push around the upwellings, due to their fixity, ultimately, the upwellings also control the locations of downwellings; regions where they impact at the surface become too hot and hence, too buoyant, to sink. Over time, this positive feedback loop creates a self-organized flow, which becomes exceptionally stable. Upwellings become more and more coherent, building up greater resistance to the background flow with time, thereby becoming anchored in their positions. We have verified that this behavior is independent of the initial condition, which is consistent with the findings of Davies (2005). These results are in agreement with the lower-vigor cases presented in Chapter 5.

As can be seen from Figure 6.3(b), the planform yielded by case 2 (i.e. a rigid upper boundary condition in an otherwise identical simulation) displays substantial variations




Figure 6.4: A series of images, displaying the thermal structure of a simulation with a free-slip upper boundary condition. Snapshots are spaced ≈ 250 Myr apart. Each one shows a radial surface just above the CMB and a hot iso-surface, representing regions of the model that are 500 K hotter than average for their depth. The scale illustrates the temperature, away from the lateral average, i.e. the thermal anomaly. The most prominent features are hot cylindrical plumes, which are robust and coherent. The figure illustrates how hot material is forced down and sideways by downwelling cold material, such that it collects at ridges. These ridges form as hot material comes together from two opposite directions. Plumes are generated where ridges meet. The tendency is for nearby plumes to join together over time and, hence, the number of plumes decreases as the simulation evolves. However, the final 'quasi' steady-state planform, which exhibits four plumes, is extremely stable.

when compared to that of case 1. Downwellings are fewer in number, although those that exist are larger and more potent. They are generated by boundary layer instabilities and, consistent with the patterns displayed in case 1, expand as they enter the lower mantle. Upwellings are more numerous, with the planform exhibiting six major plumes. When compared to surrounding mantle, they are not as hot or prominent as those in case 1, but nonetheless remain anchored in their positions (see Figure 6.5). The immobile surface, together with the increase in viscosity with depth, which combine to restrict plume movement,




Figure 6.5: As in Figure 6.4, but with a rigid upper boundary condition. Here, iso-surfaces represent regions of the model that are 350 K hotter than average for their depth. The scale illustrates the temperature, away from the lateral average, i.e. the thermal anomaly. The most prominent features are hot stationary cylindrical plumes. Once formed, these remain anchored in their positions. Indeed, excluding a few minor reorganizations, their positions are not modified following the initial transient.

is the main reason for this. Plumes cannot drift towards each other and coalesce as the simulation evolves, which accounts for the observed disparity in the number and nature of upwellings between cases 1 and 2. Another important observation is the increase in small-scale, passive upwelling activity here, when compared to case 1 (refer to Figure 6.3b). This activity, which takes on the form of convective rolls, occurs for a number of reasons. Firstly, when compared to a simulation with a free-slip surface, less heat is transferred by plumes, hence heat tends to rise in a different manner. As well as this, as the larger downwellings descend, hot material is forced upwards to compensate for the downward mass flux. However, perhaps the most important factor is that in a simulation with a free-slip surface, passive

upwellings of this nature are advected by large scale surface driven flow and although they might occur, they will be less prominent.

In examining case 3 (Figure 6.3c), we see that the imposition of plate-like surface velocities extensively restructures mantle dynamics, organizing the convective fluid by instilling the wavelength of plate geometry throughout the system. The planform is dominated by narrow, cold downwellings, which occur mainly beneath plate boundaries. However, smaller scale downwellings are visible away from plate boundaries, resulting from boundary layer instabilities. As with other cases, they widen substantially with depth and break up into large diffuse blobs in the lower mantle. This once again illustrates the importance of depth-dependent rheologies in mantle convection simulations.

The most noticeable alteration in dynamics between this and other cases is the nature of upwelling structures. The formation of active, focussed columnar upwellings is prevalent in this simulation. Their morphology is generally plume-like, however, they no longer dominate the planform, being nowhere near as coherent, or as numerous as those in cases 1 and 2. They are swept around by cold flow spreading out at the CMB and are less stationary. Indeed, by forcing subduction through the kinematic boundary condition, the balance of the system is lost. Consequently, the self-organization between upwelling and downwelling activity, clearly apparent in other cases, no longer exists. Plumes therefore cannot build up a resistance to the background flow. As with case 2, and, for the same reasons, there is evidence for passive upwelling activity throughout the mantle. These results are consistent with previous studies (Lux et al. 1979, Zhong et al. 2000, Lowman et al. 2004, Quere & Forte 2006).

The impact of the upper boundary condition is not restricted to convective structures. Integral quantities such as the *RMS* velocity, surface heat flow and the temperature profile




Figure 6.6: Temperature profile with depth for cases 1 (free-slip), 2 (rigid) and 3 (kinematic - plate motion history). Where plate motions are prescribed, mean global temperatures are greater than those of an otherwise identical system featuring a free-slip surface. Systems with a rigid surface are characterized by even greater temperatures. The temperature drop across the top boundary layer spans 42% of the temperature range in case 1, 53% in case 2 and 59% in case 3.

are also effected. Figure 6.6 is a plot of mean layer temperature against depth for all three cases. Each case varies dramatically. Where plate motions are prescribed, mean global temperatures are greater than those of an otherwise identical system featuring a free-slip surface. This is to be expected. In a free-slip simulation, active upwellings dominate the planform, transferring heat efficiently away from the CMB to the surface, where it escapes freely. However, when plate-motion is prescribed, upwelling material is swept away by background flow and often does not reach the surface, remaining insulated below the thermal boundary layer. Heat therefore has difficulty in escaping and the mantle warms up as a result. The increased Nusselt number of case 1 when compared to that of case 3 (Table 6.2), illustrating the higher surface heat flux, provides further evidence to support this conjecture.

Case	Nu_{Top}	Nu_{Bot}	V'_{RMS}
1	22.66	10.25	591.43
2	16.35	4.05	388.02
3	19.31	6.97	571.55

Table 6.2: Final Nusselt numbers and non-dimensional RMS velocities for cases 1 (free-slip), 2 (rigid) and 3 (kinematic - plate motion history). Nu_{Top} and Nu_{Bot} represent Nusselt numbers at outer and inner surfaces respectively. They differ since the models incorporate internal heating. V'_{RMS} denotes the non-dimensional global RMS velocity. Velocities are non dimensionalized by $v' = \frac{v l_0}{\kappa}$, where κ represents thermal diffusivity (calculated assuming a density, ρ , of 4.5 kg m^{-3}) and $l_0 = 2890\text{km}$.

In a simulation with a rigid surface we see that mean layer temperatures are even greater. Although plumes are prevalent in such simulations, they are not as significant as those in case 1, whilst they impinge on what is essentially a stagnant lid. With negligible advective heat transport at the surface, less heat escapes than in previous cases and, consequently, the mantle heats up. The Nusselt number displayed in Table 6.2 is in harmony with this theory. Nusselt numbers at the base of the shell are also consistent with these findings. In cases 2 and 3, i.e. those simulations with a hotter interior, the temperature gradient across the CMB is lower, hence less heat enters the mantle at its base. We note that our results do not have inverted temperature profiles like those of Monnereau & Quere (2001). This is likely due to the variation in material properties between models and the inclusion of compressibility in our simulations. Nonetheless, our general conclusions are consistent, with results, in their entirety, suggesting that the imposed plate velocity boundary condition is intermediate between two asymptotic conditions, i.e. free-slip and rigid.

6.5 Conclusions

By utilizing a modified and benchmarked version of the well-known 3-D spherical mantle-dynamics code TERRA, the influence of distinctive mechanical upper boundary conditions (free-slip, rigid and imposed plate motion) upon the internal dynamics of an otherwise

identical system has been examined. Results demonstrate, unquestionably, that the surface plays a major role in shaping the dynamics of the interior, not only influencing convective structures, but also temperature profiles and integral quantities such as *RMS* velocities. To put it simply, a change in surface conditions completely reorganizes internal behavior. Although this has been suggested by previous studies, we are now left with little doubt about the importance of the surface in mantle convection simulations. This should therefore be of major concern to the geodynamical community; the surface boundary conditions utilized in the majority of mantle convection models are not at all representative of Earth, yet, geodynamicists attempt to reconcile their models against geophysical, geochemical and geological observations on Earth. With the upper boundary playing such a fundamental role, how is it that results are at all relevant to our planet?

By failing to address the surface appropriately, geodynamical models, regardless of their sophistication, cannot be truly applicable to Earth. Although plate tectonics is a complicated mechanism, with the physics of plate generation and plate-boundary evolution far from a simple concept, the self-consistent generation of plates from mantle convection simulations must be a key focus of future research. At present, such models are very much in their infancy. However, this situation needs to be addressed if geodynamical research is to continue probing the frontiers of our geophysical understanding.

Chapter 7

Conclusions

Due to the geometrical and rheological complexities involved with geodynamical flows, numerical modeling is, without doubt, one of the best techniques available for examining such phenomena. Indeed, it is often the only feasible option. However, as has been stated many times herein, numerical models yield inaccurate results if the underlying grid cannot capture a simulation's important features. For the complex problems encountered in geodynamics, the only way to solve such flows efficiently and accurately is through multi-resolution numerical techniques. Such techniques have been developed and validated in this thesis and, subsequently, applied to a range of geophysical scenarios. Results are extremely encouraging.

The first technique introduced is grid adaptivity, a method commonly employed within the field of computational engineering. The technique, based upon unstructured discretizations, adapts the mesh automatically around regions of high solution gradient, yielding enhanced resolution of the associated flow features. Here, it is utilized in two-dimensional (2-D) convection dominated problems, using the well established CONMAN code as a foundation. When applied to purely thermal convection, the results obtained are extremely positive.

The error indicator, which is used to dictate the required grid resolution, is shown to be reliable, while the adaptive procedure, where the domain is completely remeshed, is shown to be robust. Predictions for heat transfer agree well with benchmark solutions, suggesting that the technique is valid and accurate. The method has also been applied to a more challenging thermo-chemical benchmark problem, employing a grid based method to solve for the compositional field. Such methods suffer from numerical diffusion, which is predominantly caused by insufficient resolution, a factor naturally addressed by grid adaptivity. Two hypotheses are tested:

1. The greater resolution made available through the adaptive methodologies will reduce artificial diffusion.
2. This reduced diffusion, in turn, will see grid based methods yielding results that are consistent with those achieved using other methods (i.e. tracer particle and marker chain).

As expected, results show that adapted grids yield large improvements over regular, uniform grids, generating less diffusive results whilst employing fewer degrees of freedom. However, results also demonstrate that grid based methods, even when coupled with the adaptive strategies, are not competitive with the other methods currently employed for the tracking of compositional heterogeneities. Regardless of this, the benefits of the adaptive strategies are clear, significantly reducing the number of degrees of freedom required to accurately solve a given problem. In addition:

1. The unstructured nature of the technique is suitable for discretizing the complex geometries encountered on Earth, e.g. subduction zones.
2. Nodes automatically cluster around zones of high solution gradient, without the need for complicated *a priori* mesh generation.

3. The reduction in the number of degrees of freedom leads to a decrease in computational processing time and memory use (both in terms of disk space and RAM), meaning that complex problems can be solved more efficiently.

Clearly, there are many benefits for geodynamical modeling, as is illustrated in Chapter 3, where the technique is applied to 2-D models of mid-ocean ridge (MOR) and subduction zone (SZ) magmatism. The conclusions from Chapter 2 are reinforced, with the proposed methodology shown to work remarkably well. The method refines the locations of thermal boundary layers wherever they are strong, at the ridge itself and along Earth's surface (MOR), and in the mantle wedge, along the margins of the descending plate and at Earth's surface (SZ). The adapted grids maintain good solution accuracy and, through a series of remeshings, display the ability to gradually improve solution quality, without significantly increasing the total number of unknowns at each stage. The advocated methods are once again shown to be computationally highly efficient. To date, successful goal-orientated/error-guided grid adaptation techniques have not been utilized within the field of geodynamics. The work included in these early chapters is therefore the first geodynamical application of such methods.

The second part of this thesis examines three-dimensional (3-D) problems, where multi-resolution is introduced through a different technique. Existing 3-D spherical mantle dynamics codes have shown their value on countless occasions, yielding otherwise unattainable information on the temporal and spatial evolution of mantle flow under a range of rheologies and boundary conditions. However, due to the sheer scale of the problems under study and despite the massively parallel computational architectures available to us today, with existing codes, at least in their current format, it is often impossible to resolve thermal and chemical boundary layers at the dynamical regime of Earth's mantle. The multi-resolution

scheme developed herein, employing the pre-existing 3-D spherical mantle dynamics code TERRA as a basis, provides a means to address these issues.

The scheme, based around the application of a multigrid solver on non-uniform structured grids, is shown to yield similar advantages to those observed with unstructured grids, being computationally highly efficient. However, unlike unstructured techniques, the advocated method negates the need to reformulate large sections of code, whilst also being conceptually simple and straightforward to implement. In the work covered herein, it is utilized to refine and solve on a radially non-uniform grid. Essentially, such a structure maintains the key benefits of TERRA's current configuration, whilst also overcoming many of its limitations. The approach is shown to be highly resourceful in terms of RAM, meaning that one can attempt calculations that would otherwise be impractical. As an added benefit, the solution algorithm reduces the CPU-time needed to solve a given problem. Validation tests illustrate that the technique is accurate and robust.

The method is applied to cutting-edge 3-D spherical geophysical models, with two different processes investigated. Firstly, the nature of upwelling structures is examined in high vigor, global mantle convection models (Chapter 5). The modified code allows one to efficiently resolve thermal boundary layers at the dynamical regime of Earth's mantle - the simulations presented achieve Rayleigh numbers of greater than 1×10^9 , the highest attained, to date, in global 3-D spherical geometry. At these vigors, models yield planforms characterized by both long-lived, steady plumes and mobile, ephemeral plumes. Such variable plume characteristics are consistent with hot-spot observations on Earth. In addition, those plumes that do remain steady vary in intensity throughout the simulation. We argue that such behavior is caused by an irregular supply of cold material to the core-mantle boundary region. Since subduction on Earth is a time-dependent and asymmetric process, these

results suggest that subducting slabs might be partially responsible for the episodic plume magmatism recorded in large flood basalt provinces.

The simulations examined imply that the surface boundary condition must be addressed more appropriately if we are to reconcile our models against geophysical, geochemical and geological observations on Earth. This point is examined further in Chapter 6, where the influence of the surface upon internal dynamics is investigated. Three different mechanical boundary conditions are modeled (free-slip, rigid and kinematic with imposed plate motion history), in otherwise identical systems. Results demonstrate that the surface strongly dictates the nature of flow within the interior, manifesting itself not only in convective structures, but also in thermal profiles, Nusselt numbers and velocity patterns. Since the majority of geodynamical simulations incorporate a surface condition that is not at all representative of Earth, this is a worrying, yet important conclusion. By failing to address the surface appropriately, geodynamical models, regardless of their sophistication, cannot be truly applicable to Earth.

In summary, the two major goals outlined in the introduction to this thesis have been accomplished. The first, to develop computationally efficient, multi-resolution numerical techniques, has been achieved in both 2- and 3-D. The second, to utilize these advanced models in cutting-edge geophysical simulations, has also been achieved, with viscosity structure, Rayleigh number and the surface boundary condition shown to be extremely important in mantle convection simulations. Indeed, in addition to chemical heterogeneities, results suggest that these features have more of an influence upon mantle dynamics than any other physical mechanism.

The methods developed have undoubtedly allowed us to solve problems for which solutions were previously unattainable and the tools are now in place for further investigations. Bear-

ing in mind the rapid advances being made in seismology, geodesy, geochemistry, mineral physics, not to forget computational hardware, exciting times lie ahead in geodynamical research. Novel techniques like those introduced here, should ensure that best use is made of computational resources, both now and in the future.

Bibliography

- Albarede, F. & van der Hilst, R. D. (2002), 'Zoned mantle convection', *Phil. Trans. Roy. Soc. Lon. A* **360**, 2569–2592.
- Albers, M. & Christensen, U. R. (2001), 'Channeling of plume flow beneath mid-ocean ridges', *Earth Planet. Sci. Lett.* **187**, 207–220.
- Alcouffe, R. E., Brandt, A., Dendy, J. E. & Painter, J. W. (1981), 'The multigrid method for the diffusion equation with strongly discontinuous coefficients', *SIAM J. Sci. Stat. Comp.* **2**, 430–454.
- Alfe, D., Gillan, M. J. & Price, G. D. (2000), 'Constraints on the composition of the Earth's core from ab initio calculations', *Nature* **405**, 172–175.
- Allegre, C. J., Brevart, O., Dupre, B. & Minster, J. F. (1980), 'Isotopic and chemical effects produced in a continuously differentiating convecting Earth mantle', *Phil. Trans. Roy. Soc. Lon. A* **297**, 447–477.
- Allegre, C. J., Hofmann, A. & O'Nions, K. (1996), 'The Argon constraints on mantle structure', *Geophys. Res. Lett.* **23**, 3555–3557.
- Anderson, D. L. (1998), 'The scales of mantle convection', *Tectonophysics* **284**, 1–17.
- Anderson, D. L. (2000), 'The thermal state of the upper mantle: no role for mantle plumes', *Geophys. Res. Lett.* **27**, 3623–3626.
- Anderson, D. L. (2001), 'Top-down tectonics?', *Science* **294**, 57–61.
- Andrews, D. J. & Sleep, N. H. (1974), 'Numerical modeling of tectonic flow behind island arcs', *Geophys. J. Roy. Astr. Soc.* **38**, 237–251.
- Atanga, J. & Silvester, D. (1992), 'Iterative methods for stabilized mixed velocity pressure finite elements', *Int. J. Num. Meth. Fluids* **14**, 71–81.
- Babuska, I. & Rheinboldt, W. C. (1978), 'A-posteriori error estimates for the finiteelement method', *Int. J. Num. Meth. Eng.* **12**, 1597–1615.
- Bai, D. & Brandt, A. (1987), 'Local mesh refinement and multilevel techniques', *SIAM J. Sci. Stat. Comp.* **8**, 109–134.

- Baker, T. J. (1990), 'Unstructured mesh generation by a generalized Delaunay algorithm, in: Applications of Mesh Generation to Complex 3D Configurations', *AGARD Conf. Proc.* **464**, 20.1–20.10.
- Baumgardner, J. R. (1983). A three-dimensional finite element model for mantle convection, PhD thesis, UCLA.
- Baumgardner, J. R. (1985), 'Three-dimensional treatment of convective flow in the Earth's mantle', *J. Stat. Phys.* **38**, 501–511.
- Baumgardner, J. R. & Frederickson, P. O. (1985), 'Icosahedral discretization of the two-sphere', *SIAM J. Num. Anal.* **22**, 1107–1115.
- Becker, T. W. & Boschi, L. (2002), 'A comparison of tomographic and geodynamic mantle models', *Geochem. Geophys. Geosys.* **3**, doi:10.1029/2001GC000168.
- Bercovici, D. (2003), 'The generation of plate tectonics from mantle convection', *Earth Planet. Sci. Lett.* **205**, 107–121.
- Bercovici, D. & Karato, S. (2003), 'Whole-mantle convection and the transition-zone water filter', *Nature* **425**, 39–44.
- Bercovici, D. & Mahoney, J. (1994), 'Double flood basalts and plume head separation at the 660 kilometer discontinuity', *Science* **266**, 1367–1369.
- Bercovici, D., Ricard, Y. & Richards, M. A. (2000), 'The relation between mantle dynamics and plate tectonics: a primer', *Geophys. Monogr. Ser.* **121**, 5–46.
- Bercovici, D., Schubert, G. & Glatzmaier, G. A. (1989), '3-D spherical models of convection in the Earth's mantle', *Science* **244**, 950–955.
- Bijwaard, H. & Spakman, W. (1999), 'Tomographic evidence for a narrow whole mantle plume below Iceland', *Earth Planet. Sci. Lett.* **166**, 121–126.
- Birch, F. (1964), 'Density and composition of mantle and core', *J. Geophys. Res.* **69**, 4377–4388.
- Blankenbach, B., Busse, F., Christensen, U., Cserepes, L., Gunkel, D., Hansen, U., Harder, H., Jarvis, G., Koch, M., Marquart, G., Moore, D., Olson, P., Schmeling, H. & Schnaubelt, T. (1989), 'A benchmark comparison for mantle convection codes', *Geophys. J. Int.* **98**, 23–38.
- Bodri, L. & Bodri, B. (1978), 'Numerical investigation of tectonic flow in island-arcs', *Tectonophysics* **50**, 163–175.
- Boyet, M. & Carlson, R. W. (2006), 'A new geochemical model for the Earth's mantle inferred from (SM)-S-146-Nd-142 systematics', *Earth Planet. Sci. Lett.* **250**, 254–268.
- Boyet, M. & Carlson, R. W. (n.d.), 'Nd-142 evidence for early (> 4.53 ga) global differentiation of the silicate earth'.

- Brandon, A. D. & Walker, R. J. (2005), 'The debate over core-mantle interaction', *Earth Planet. Sci. Lett.* **232**, 211–225.
- Brandt, A. (1977), 'Multilevel adaptive solutions to boundary value problems', *Math. Comp.* **31**, 333–390.
- Brandt, A. (1984), *Multigrid techniques: 1984 guide with applications to fluid dynamics*, GMD.
- Briggs, W. L., Henson, V. E. & McCormick, S. F. (2000), *A multigrid tutorial*, 2nd edn, SIAM, United States.
- Buck, W. R., Lavier, L. L. & Poliakov, A. N. B. (2005), 'Modes of faulting at mid-ocean ridges', *Nature* **434**, 719–723.
- Bullard, E. C. (1939), 'Heat flow in South Africa', *Proc. Roy. Soc. London, Ser. A.* **173**, 474–502.
- Bunge, H. P. & Baumgardner, J. R. (1995), 'Mantle convection modeling on parallel virtual machines', *Comp. Phys.* **9**, 207–215.
- Bunge, H. P. & Davies, J. H. (2001), 'Tomographic images of a mantle circulation model', *Geophys. Res. Lett.* **28**, 77–80.
- Bunge, H. P. & Grand, S. P. (2000), 'Mesozoic plate-motion history below the northeast Pacific Ocean from seismic images of the subducted Farallon slab', *Nature* **405**, 337–340.
- Bunge, H. P., Hagelberg, C. R. & Travis, B. J. (2003), 'Mantle circulation models with variational data assimilation: inferring past mantle flow and structure from plate motion histories and seismic tomography', *Geophys. J. Int.* **152**, 280–301.
- Bunge, H. P., Richards, M. A. & Baumgardner, J. R. (1996), 'Effect of depth-dependent viscosity on the planform of mantle convection', *Nature* **379**, 436–438.
- Bunge, H. P., Richards, M. A. & Baumgardner, J. R. (1997), 'A sensitivity study of three-dimensional spherical mantle convection at 10^8 Rayleigh number: Effects of depth-dependent viscosity, heating mode, and an endothermic phase change', *J. Geophys. Res.* **102**, 11991–12007.
- Bunge, H. P., Richards, M. A. & Baumgardner, J. R. (2002), 'Mantle-circulation models with sequential data assimilation: inferring present-day mantle structure from plate-motion histories', *Phil. Trans. Roy. Soc. Lon. A* **360**, 2545–2567.
- Bunge, H. P., Richards, M. A., Lithgow-Bertelloni, C., Baumgardner, J. R., Grand, S. P. & Ramanowicz, B. A. (1998), 'Time scales and heterogeneous structure in geodynamic Earth models', *Science* **280**, 91–95.
- Busse, F. H. & Riahi, N. (1982), 'Patterns of convection in spherical shells .2.', *J. Fluid Mech.* **123**, 283–301.

- Carlsaw, H. S. & Jaeger, J. C. (1959), *Conduction of heat in solids*, 2nd edn, Oxford University Press, New York.
- Chambers, J. E. (2001), 'Making more terrestrial planets', *ICARUS* **152**, 205–224.
- Chase, C. G. (1979), 'Subduction, the geoid, and lower mantle convection', *Nature* **282**, 464–468.
- Chase, C. G. (1985), 'The geological significance of the geoid', *Ann. Rev. Earth Planet. Sci.* **13**, 97–117.
- Choblet, G., Cadek, O., Couturier, F. & Dumoulin, C. (2007), 'OEDIPUS: a new tool to study the dynamics of planetary interiors', *Geophys. J. Int.* **170**, 9–30.
- Chopelas, A. & Boehler, R. (1992), 'Thermal expansivity in the lower mantle', *Geophys. Res. Lett.* **19**, 1983–1986.
- Christensen, U. (2001), 'Geodynamic models of deep subduction', *Phys. Earth Planet. Int.* **127**, 25–34.
- Christensen, U. R. (1989), 'The heat transport by convection rolls with free boundaries at high Rayleigh number', *Geophys. Astrophys. Fluid Dyn.* **46**, 93–103.
- Coffin, M. F., Pringle, M. S., Duncan, R. A., Gladczenco, T. P., Storey, M., Muller, R. D. & Gahagan, L. A. (2002), 'Kergulen hotspot magma output since 130 Ma', *J. Petrol.* **43**, 1121–1139.
- Conrad, C. P. & Hager, B. H. (1999), 'The thermal evolution of an Earth with strong subduction zones', *Geophys. Res. Lett.* **26**, 3041–3044.
- Courtillot, V., Davaille, A., Besse, J. & Stock, J. (2003), 'Three distinct types of hotspots in the Earth's mantle', *Earth Planet. Sci. Lett.* **205**, 295–308.
- Cuthill, E. H. & McKee, J. (1969), 'Reducing the bandwidth of sparse symmetric matrices. Paper presented at ACM 24th National Conference, Assoc. for Comput. Mech., New York'.
- Davaille, A. (1999), 'Simultaneous generation of hotspots and superswells by convection in a heterogeneous planetary mantle', *Nature* **402**, 756–760.
- Davaille, A., Girard, F. & Bars, M. L. (2002), 'How to anchor hotspots in a convecting mantle?', *Earth Planet. Sci. Lett.* **203**, 621–634.
- Davies, D. R., Davies, J. H., Hassan, O., Morgan, K. & Nithiarasu, P. (2007), 'Investigations into the applicability of adaptive finite element methods to two-dimensional infinite Prandtl number thermal and thermochemical convection', *Geochem. Geophys. Geosys.* **8**, Q05010, doi:10.1029/2006GC001470.
- Davies, G. F. (1988), 'Ocean bathymetry and mantle convection .1. large-scale flow and hotspots', *J. Geophys. Res.* **93**, 10467–10480.

- Davies, G. F. (1989), 'Mantle convection model with a dynamic plate - topography, heat-flow and gravity-anomalies', *Geophys. J. Int.* **98**, 461–464.
- Davies, G. F. (1992), 'Temporal variation of the Hawaiian plume flux', *Earth Planet. Sci. Lett.* **113**, 277–286.
- Davies, G. F. (1999), *Dynamic Earth: Plates, Plumes and Mantle Convection*, Cambridge University Press, Cambridge.
- Davies, G. F. & Richards, M. A. (1992), 'Mantle Convection', *J. Geol.* **100**, 151–206.
- Davies, J. H. (2005), 'Steady plumes produced by downwellings in Earth-like vigor spherical whole mantle convection models', *Geochem. Geophys. Geosys.* **6**, Q12001, doi:10.1029/2005GC001042.
- Davies, J. H. & Bunge, H. P. (2006), 'Are splash plumes the origin of minor hotspots?', *Geology* **34**, 349–352.
- Davies, J. H. & Stevenson, D. J. (1992), 'Physical model of source region of subduction zone volcanics', *J. Geophys. Res.* **97**, 2037–2070.
- Demkowicz, L., Oden, J. T. & Strouboulis, T. (1984), 'Adaptive finite-elements for flow problems with moving boundaries - variational principles and a posteriori estimates', *Comp. Meth. Appl. Mech.* **46**, 217–251.
- Depaolo, D. J. & Wasserburg, G. J. (1976), 'Inferences about magma sources and mantle structure from variations of Nd-143-Nd-144', *Geophys. Res. Lett.* **3**, 743–746.
- Dziewonski, A. M. & Anderson, D. L. (1981), 'Preliminary Reference Earth Model', *Phys. Earth Planet. Int.* **25**, 297–356.
- Dziewonski, A. M., Hales, A. L. & Lapwood, E. R. (1975), 'Parametrically simple Earth models consistent with geophysical data', *Phys. Earth Planet. Int.* **10**, 12–48.
- Ekevid, T., Kettil, P. & Wiberg, N. E. (2004), 'Adaptive multigrid for finite element computations in plasticity', *Comp. Struct.* **82**, 2413–2424.
- El Hachemi, M., Hassan, O., Morgan, K., Rowse, D. P. & Weatherill, N. P. (2003), 'Hybrid methods for electromagnetic scattering simulations on overlapping grids', *Comm. Num. Meth. Eng.* **19**, 749–760.
- Farnetani, D. G. & Richards, M. A. (1995), 'Thermal entrainment and melting in mantle plumes', *Earth Planet. Sci. Lett.* **136**, 251–267.
- Finn, C. A., Muller, R. D. & Panter, K. S. (2005), 'A Cenozoic diffuse alkaline magmatic province (DAMP) in the southwest Pacific without rift or plume origin', *Geochem. Geophys. Geosys.* **6**, Q02005, doi:10.1029/2004GC000723.
- Formaggia, L., Peraire, J., Morgan, K. & Peiro, J. (1988), 'Implementation of a 3-D explicit Euler solver on a CRAY computer, in: Proc. 4th Int. Symposium on Science and Engineering on CRAY Supercomputers, Minneapolis', pp. 45–65.

- Forsyth, D. & Uyeda, S. (1975), 'On the relative importance of the driving forces of plate motion', *Geophys. J. Roy. Astr. Soc.* **43**, 163–200.
- Forte, A. M., Mitrovica, J. M. & Espeset, A. (2002), 'Geodynamic and seismic constraints on the thermochemical structure and dynamics of convection in the deep mantle', *Phil. Trans. Roy. Soc. Lon. A* **360**, 2521–2543.
- Fortin, M., Vallet, M. G., Dompierre, J., Bourgault, J. & Habashi, W. G. (1996), 'Anisotropic mesh adaptation: Theory, validation and applications', *Comp. Fluid Dyn.* **96**, 174–180.
- Foulger, G. R., Pritchard, M. J. & Julian, B. R. (2001), 'Seismic tomography shows that upwelling beneath Iceland is confined to the upper mantle', *Geophys. J. Int.* **146**, 504–530.
- Fullsack, P. (1995), 'An arbitrary Lagrangian-Eulerian formulation for creeping flows and its application in tectonic models', *Geophys. J. Int.* **120**, 1–23.
- Gable, C. W., Oconnell, R. J. & Travis, B. J. (1991), 'Convection in 3-D with surface plates - generation of toroidal flow', *J. Geophys. Res.* **96**, 8391–8405.
- George, A. J. (1971), Computer implementation of the finite element method, PhD thesis, STAN-CS-71-208, Stanford Univ., Stanford, California, USA.
- Gessmann, C. K., Wood, B. J., Rubie, D. C. & Kilburn, M. R. (2001), 'Solubility of silicon in liquid metal at high pressure: implications for the composition of the Earth's core', *Earth Planet. Sci. Lett.* **184**, 367–376.
- Goes, S., Capitanio, F. A. & Morra, G. (2008), 'Evidence of lower-mantle slab penetration phases in plate motions', *Nature* **451**, 981–984.
- Gonnermann, H. M. & Mukhopadhyay, S. (2007), 'Non-equilibrium degassing and a primordial source for helium in ocean-island volcanism', *Nature* **449**, 1037–1040.
- Grand, S. P. (2002), 'Mantle shear-wave tomography and the fate of subducted slabs', *Phil. Trans. Roy. Soc. Lon. A* **360**, 2475–2491.
- Grand, S. P., van der Hilst, R. & Widiyantoro, S. (1997), 'Global seismic tomography: a snapshot of convection in the Earth', *Geol. Soc. Am. Today* **7**, 1–7.
- Gurnis, M. & Davies, G. F. (1986), 'Mixing in numerical-models of mantle convection incorporating plate kinematics', *J. Geophys. Res.* **91**, 6375–6395.
- Gurnis, M. & Hager, B. H. (1988), 'Controls on the structure of subducted slabs', *Nature* **335**, 317–321.
- Gutenberg, B. (1913), 'Über die Konstitution der Erdinnern, erschlossen aus Erdbebenbeobachtungen', *Phys. Zeit.* **14**, 1217.
- Hager, B. H. (1984), 'Subducted slabs and the geoid - constraints on mantle rheology and flow', *J. Geophys. Res.* **89**, 6003–6019.

- Hager, B. H. & O'Connell, R. J. (1981), 'A simple global model of plate dynamics and mantle convection', *J. Geophys. Res.* **86**, 4843–4867.
- Han, L. J. & Gurnis, M. (1999), 'How valid are dynamic models of subduction and convection when plate motions are prescribed?', *Phys. Earth Planet. Int.* **110**, 235–246.
- Hansen, U., Yuen, D. A., Kroening, S. E. & Larsen, T. B. (1993), 'Dynamic consequences of depth-dependent thermal expansivity and viscosity on mantle circulations and thermal structure', *Phys. Earth Planet. Int.* **77**, 205–223.
- Harder, H. & Christensen, U. R. (1996), 'A one-plume model of martian mantle convection', *Nature* **380**, 507–509.
- Hassan, O. & Probert, E. J. (1999), 'Grid control and adaptation. in: Handbook of Grid Generation', pp. (35)1–(35)29.
- Hassan, O., Probert, E. J., Morgan, K. & Peraire, J. (1995), 'Mesh generation and adaptivity for the solution of compressible viscous high-speed flows', *Int. J. Num. Meth. Eng.* **38**, 1123–1148.
- Hewitt, J. M., McKenzie, D. P. & Weiss, N. O. (1980), 'Large aspect ratio cells in two-dimensional thermal convection', *Earth Planet. Sci. Lett.* **51**, 370–380.
- Hieronymus, C. F. & Bercovici, D. (1999), 'Discrete alternating hotspot islands formed by interaction of magma transport and lithospheric flexure', *Nature* **397**, 604–607.
- Hoffman, A. W. (1997), 'Mantle chemistry: the message from oceanic volcanism.', *Nature* **385**, 219–229.
- Hughes, T. J. R. (1980), 'Implicit-Explicit finite element techniques for symmetric and non-symmetric systems, *Proc. First. Int. Conf. Num. Meth. for non-linear problems*, Swansea, UK'.
- Hughes, T. J. R. (2000), *The Finite Element Method: Linear Static and Dynamic Finite Element Analysis*, Dover Publications.
- Hughes, T. J. R. & Brooks, A. (1979), 'A multi-dimensional upwind scheme with no crosswind diffusion. in: Finite element methods for convection dominated flows', **34**, 19–35.
- Hughes, T. J. R., Liu, W. K. & Brooks, A. (1979), 'Finite element analysis of incompressible viscous flows by the penalty function formulation', *J. Comp. Phys.* **30**, 1–60.
- Iwamori, H. (1998), 'Transportation of H₂O and melting in subduction zones', *Earth Planet. Sci. Lett.* **160**, 65–80.
- Jarvis, G. T. (1984), 'Time-dependent convection in the Earth's mantle', *Phys. Earth Planet. Int.* **36**, 305–327.
- Jeffreys, H. (1926), 'The rigidity of Earth's central core', *Mon. Not. Roy. Astron. Soc. Geophys. Suppl.* **1**, 371–383.

- Jellinek, A. M. & Manga, M. (2002), 'The influence of a chemical boundary layer on the fixity, spacing and lifetime of mantle plumes', *Nature* **418**, 760–763.
- Jessop, A. M., Hobart, M. & Sclater, J. G. (1976), 'World Wide Compilation of Heat Flow Data', *Geo. Ser.* **5**.
- Jurdy, D. M. & Gordon, R. G. (1984), 'Global plate motions relative to the hot spots 64 to 56 Ma', *J. Geophys. Res.* **89**, 9927–9936.
- Kameyama, M., Kageyama, A. & Sato, T. (2005), 'Multigrid iterative algorithm using pseudo-compressibility for three-dimensional mantle convection with strongly variable viscosity', *J. Comp. Phys.* **206**, 162–181.
- Kellogg, L. H., Hager, B. H. & van der Hilst, R. D. (1999), 'Compositional stratification in the deep mantle', *Science* **283**, 1881–1884.
- Kennett, B. L. N. & Engdahl, E. R. (1991), 'Traveltimes for global earthquake location and phase identification', *Geophys. J. Int.* **105**, 429–465.
- Kennett, B. L. N. & Gorbатов, A. (2004), 'Seismic heterogeneity in the mantle - strong shear wave signature of slabs from joint tomography', *Phys. Earth Planet. Int.* **146**, 87–100.
- Kincaid, C. & Sacks, I. S. (1997), 'Thermal and dynamical evolution of the upper mantle in subduction zones', *J. Geophys. Res.* **102**, 12295–12315.
- King, S. D. & Hager, B. H. (1994), 'Subducted slabs and the geoid .1. Numerical experiments with temperature-dependent viscosity', *J. Geophys. Res.* **99**, 19843–19852.
- King, S. D., Lowman, J. P. & Gable, C. W. (2002), 'Episodic tectonic plate reorganizations driven by mantle convection', *Earth Planet. Sci. Lett.* **203**, 83–91.
- King, S. D., Raefsky, A. & Hager, B. H. (1990), 'CONMAN - vectorizing a finite-element code for incompressible 2-D convection in the Earth's mantle', *Phys. Earth Planet. Int.* **59**, 195–207.
- Korenaga, J. (2005), 'Firm mantle plumes and the nature of the core-mantle boundary region', *Earth Planet. Sci. Lett.* **232**, 29–37.
- Kuhn, D. & Dahm, T. (2004), 'Simulation of magma ascent by dykes in the mantle beneath mid-ocean ridges', *J. Geodyn.* **38**, 147–159.
- Lehmann, I. (1936), 'P'', *Trav. Sci., Sect. Seis. U. G. G. I. (Toulouse)* **14**, 3–31.
- Lenardic, A. & Kaula, W. M. (1993), 'A numerical treatment of geodynamic viscous-flow problems involving the advection of material interfaces', *J. Geophys. Res.* **98**, 8243–8260.
- Lewis, R., Nithiarasu, P. & Seetharamu, K. (2004), *Fundamentals of the Finite Element Method for Heat and Fluid Flow*, Wiley, New York.
- Li, J. & Agee, C. B. (2001), 'Element partitioning constraints on the light element composition of the Earth's core', *Geophys. Res. Lett.* **28**, 81–84.

- Lin, S.-C. & van Keken, P. E. (2005), 'Multiple volcanic episodes of flood basalts caused by thermochemical mantle plumes', *Nature* **436**, 250–252.
- Lissauer, J. J. (1993), 'Planet formation', *Ann. Rev. Astron. Astrophys.* **31**, 129–174.
- Lithgow-Bertelloni, C. & Richards, M. A. (1998), 'The dynamics of Cenozoic and Mesozoic plate motions', *Rev. Geophys.* **36**, 27–78.
- Lo, S. H. (1985), 'A new mesh generation scheme for arbitrary planar domains', *Int. J. Num. Meth. Eng.* **21**, 1403–1426.
- Lohner, R. (1995), 'Mesh adaptation in fluid-mechanics', *Eng. Fract. Mech.* **50**, 819–847.
- Lohner, R. (2001), *Applied CFD techniques*, Wiley, New York.
- Lohner, R., Morgan, K. & Zienkiewicz, O. C. (1985), 'An adaptive finite element method for high-speed compressible flow', *Lect. Notes Phys.* **218**, 388–392.
- Lohner, R., Morgan, K. & Zienkiewicz, O. C. (1986), 'Adaptive grid refinement for the Euler and compressible Navier-Stokes equations. in: Accuracy Estimates and Adaptive Refinement in Finite Element Computations.', pp. 281–296.
- Lopez, S. & Casciaro, R. (1997), 'Algorithmic aspects of adaptive multigrid finite element analysis', *Int. J. Num. Meth. Eng.* **140**, 919–936.
- Lopez, S., Fortino, S. & Casciaro, R. (1998), 'An adaptive multigrid solver for plate vibration and buckling problems', *Comp. Struct.* **69**, 625–637.
- Lowman, J. P., King, S. D. & Gable, C. W. (2001), 'The influence of tectonic plates on mantle convection patterns, temperature and heat flow', *Geophys. Res. Lett.* **146**, 619–636.
- Lowman, J. P., King, S. D. & Gable, C. W. (2004), 'Steady plumes in viscously stratified, vigorously convecting, three-dimensional numerical mantle convection models with mobile plates', *Geochem. Geophys. Geosys.* **5**, Q01L01, doi:10.1029/2003GC000583.
- Lux, R. A., Davies, G. F. & Thomas, J. H. (1979), 'Moving lithospheric plates and mantle convection', *Geophys. J. Roy. Astr. Soc.* **58**, 209–228.
- Masters, T. G. & Shearer, P. M. (1990), 'Summary of seismological constraints on the structure of Earth's core', *J. Geophys. Res.* **95**, 21691–21695.
- Matyska, C. & Yuen, D. A. (2001), 'Are mantle plumes adiabatic?', *Earth Planet. Sci. Lett.* **189**, 165–176.
- Mayne, D. A., Usmani, A. S. & Crapper, M. (2000), 'h-adaptive finite element solution of high Rayleigh number thermally driven cavity problems', *Int. J. Num. Meth. Heat Fluid Flow* **10**, 598–615.
- McKenzie, D. P. (1969), 'Speculations on the consequences and causes of plate motions', *Geophys. J. Roy. Astr. Soc.* **18**, 1–32.

- McKenzie, D. P., Roberts, J. M. & Weiss, N. O. (1973), 'Numerical models of convection in the Earth's mantle', *Tectonophysics* **19**, 89–103.
- McKenzie, D. P., Roberts, J. M. & Weiss, N. O. (1974), 'Convection in Earth's mantle - towards a numerical simulation', *J. Fluid Mech.* **62**, 465–538.
- McNamara, A. K. & Zhong, S. J. (2005), 'Thermochemical structures within a spherical mantle: Superplumes or piles?', *J. Geophys. Res.* **109**, B07402.
- Milne, J. (1906), 'Bakerian Lecture - recent advances in seismology', *Proc. Roy. Soc. A* **77**, 365–376.
- Mitrovica, J. X. (1996), 'Haskell [1935] revisited', *J. Geophys. Res.* **101**, 555–569.
- Mitrovica, J. X. & Forte, A. M. (2004), 'A new inference of mantle viscosity based upon joint inversion of convection and glacial isostatic adjustment data', *Earth Planet. Sci. Lett.* **225**, 177–189.
- Mitrovica, J. X. & Jarvis, G. T. (1987), 'A numerical study of thermal convection between rigid horizontal boundaries', *Geophys. Astrophys. Fluid Dyn.* **38**, 193–223.
- Mitrovica, J. X. & Peltier, W. R. (1995), 'Constraints on mantle viscosity based upon the inversion of post-glacial rebound data from the Hudson Bay region', *Geophys. J. Int.* **122**, 353–377.
- Molnar, P. & Stock, J. (1987), 'Relative motions of hotspots in the Pacific, Atlantic and Indian Oceans since Late Cretaceous time', *Nature* **327**, 587–591.
- Monnereau, M. & Quere, S. (2001), 'Spherical shell models of mantle convection with tectonic plates', *Earth Planet. Sci. Lett.* **184**, 575–587.
- Montelli, R., Nolet, G., Dahlen, F. A., Masters, G., Engdahl, E. R. & Hung, S. H. (2004), 'Finite-frequency tomography reveals a variety of plumes in the mantle', *Science* **303**, 338–343.
- Moresi, L. N. & Solomatov, V. S. (1995), 'Numerical investigation of 2-D convection with extremely large viscosity variations', *Phys. Fluids* **7**, 2154–2162.
- Morgan, W. J. (1972), 'Plate motions and deep mantle convection', *Mem. Geol. Soc. Am.* **132**, 7–22.
- Morgan, W. J. (1983), 'Hotspot tracks and the early rifting of the Atlantic', *Tectonophysics* **94**, 123–129.
- Murakami, M., Hirose, K., Kawamura, K., Sata, N. & Ohishi, Y. (2004), 'Phase transition of MgSiO₃ perovskite in the deep lower mantle', *Lithos.* **73**, 855–858.
- Murnaghan, F. D. (1951), *Finite deformation on an elastic solid*, Wiley, New York.

- Nakagawa, T. & Tackley, P. J. (2006), 'Three-dimensional structures and dynamics in the deep mantle: Effects of post-perovskite phase change and deep mantle layering', *Geophys. Res. Lett.* **33**, L12S11, doi:10.1029/2006GL025719.
- Nettelfield, D. & Lowman, J. P. (2007), 'The influence of plate-like surface motion on upwelling dynamics in numerical mantle convection models', *Phys. Earth Planet. Int.* **161**, 184–201.
- Ni, S. D., Tan, E., Gurnis, M. & Helmberger, D. (2002), 'Sharp sides to the African super-plume', *Science* **296**, 1850–1852.
- Nielson, G. M. (1979), 'Side-vertex method for interpolation in triangles', *J. Approx. Theory* **25**, 318–336.
- Nithiarasu, P. (2000), 'An adaptive finite element procedure for solidification problems', *Heat Mass Trans.* **36**, 223–229.
- Nithiarasu, P. & Zienkiewicz, O. C. (2000), 'Adaptive mesh generation for fluid mechanics problems', *Int. J. Num. Meth. Eng.* **47**, 629–662.
- Norton, I. O. (1995), 'Plate motions in the North Pacific - the 43Ma Movement', *Tectonics* **14**, 1080–1094.
- Oganov, A. R. & Ono, S. (2004), 'Theoretical and experimental evidence for a post-perovskite phase of MgSiO₃ in Earth's D" layer', *Nature* **430**, 445–448.
- Oldham, D. N. (2004), On the possibility of layered mantle convection; Numerical simulations in spherical geometry, PhD thesis, Cardiff University, UK.
- Oldham, D. N. & Davies, J. H. (2004), 'Numerical investigation of layered convection in a three-dimensional shell with application to planetary mantles', *Geochem. Geophys. Geosys.* **5**, Q12C04, doi:10.1029/2003GC000603.
- Oldham, R. D. (1906), 'The constitution of the Earth as revealed by earthquakes', *Quart. J. Geol. Soc.* **62**, 456–475.
- Olson, P. (1987), 'A comparison of heat transfer laws for mantle convection at very high Rayleigh numbers', *Phys. Earth Planet. Int.* **48**, 153–160.
- O'Nions, R. K. & Oxburgh, E. R. (1983), 'Heat and helium in the Earth', *Nature* **306**, 429–431.
- Parman, S. W. (2007), 'Helium isotopic evidence for episodic mantle melting and crustal growth', *Nature* **446**, 900–903.
- Peacock, S. M. (1996), 'Thermal and petrologic structure of subduction zones. in: Subduction Top to Bottom', pp. 1–32.
- Pearson, D. G., Parman, S. W. & Nowell, G. M. (2007), 'A link between large mantle melting events and continent growth seen in Osmium isotopes', *Nature* **449**, 202–205.

- Pelletier, D. & Ilinca, F. (1995), 'Adaptive finite element method for mixed convection', *J. Thermo. Heat Trans.* **9**, 708–714.
- Peraire, J., Morgan, K. & Peiro, J. (1990), 'Unstructured finite element mesh generation and adaptive procedures for CFD, *AGARD Conf. Proc.*', **464**, 18.1–18.12.
- Peraire, J., Vahdati, M., Morgan, K. & Zienkiewicz, O. C. (1987), 'Adaptive remeshing for compressible flow calculations', *J. Comp. Phys.* **72**, 449–466.
- Poirier, J. P. (1994), 'Light-elements in Earth's outer core - a critical review', *Phys. Earth Planet. Int.* **85**, 319–337.
- Pollack, H. N., Hurter, S. J. & Johnson, J. R. (1993), 'Heat flow from the Earth's interior - analysis of the global data set', *Rev. Geoph.* **31**, 267–280.
- Quere, S. & Forte, A. M. (2006), 'Influence of past and present-day plate motions on spherical models of mantle convection: implications for mantle plumes and hotspots', *Geophys. J. Int.* **165**, 1041–1057.
- Ramage, A. & Wathen, A. J. (1994), 'Iterative solution techniques for the Stokes and Navier-Stokes equations', *Int. J. Num. Meth. Fluids* **19**, 67–83.
- Ranalli, G. (2001), 'Mantle rheology: Radial and lateral viscosity variations inferred from microphysical creep laws', *J. Geodyn* **32**, 425–444.
- Ratcliff, J. T., Bercovici, D., Schubert, G. & Kroenke, L. W. (1998), 'Mantle plume heads and the initiation of plate tectonic reorganizations', *Earth Planet. Sci. Lett.* **156**, 195–207.
- Ratcliff, J. T., Schubert, G. & Zebib, A. (1996), 'Steady tetrahedral and cubic patterns of spherical shell convection with temperature-dependent viscosity', *J. Geophys. Res.* **101**, 25473–25484.
- Ricard, Y., Richards, M., Lithgow-Bertelloni, C. & Lestunff, Y. (1993), 'A geodynamic model of mantle density heterogeneity', *J. Geophys. Res.* **98**, 21895–21909.
- Ricard, Y., Sabadini, R. & Spada, G. (1992), 'Isostatic deformations and polar wander induced by redistribution of mass within the Earth', *J. Geophys. Res.* **97**, 14223–14236.
- Richards, M. A., Duncan, R. A. & Courtillot, V. (1989), 'Flood basalts and hot-spot tracks: plume heads and tails', *Science* **246**, 103–107.
- Richards, M. A. & Engebretson, D. C. (1992), 'Large-scale mantle convection and the history of subduction', *Nature* **355**, 437–440.
- Richards, M. A. & Hager, B. H. (1984), 'Geoid anomalies in a dynamic Earth', *J. Geophys. Res.* **89**, 5987–6002.
- Rivara, M. C. (1984), 'Algorithms for refining triangular grids suitable for adaptive and multigrid techniques', *Int. J. Num. Meth. Eng.* **20**, 745–756.

- Roberts, G. O. (1979), 'Fast viscous Benard convection', *Geophys. Astrophys. Fluid Dyn.* **12**, 235–272.
- Rude, U. (1993), 'Fully adaptive multigrid methods', *SIAM J. Num. Anal.* **30**, 230–248.
- Rutherford, E. (1907), 'Some cosmical aspects of radioactivity', *J. Roy. Astr. Soc. Canada* **62**, 145–165.
- Safronov, V. (1969), 'Evolution of the Protoplanetary Cloud and Formation of the Earth and Planets. Moscow: Nauka. English translation, NASA TTF-677'.
- Schmerr, N. & Garnero, E. J. (2007), 'Upper mantle discontinuity topography from thermal and chemical heterogeneity', *Science* **318**, 623–626.
- Schubert, G. & Anderson, C. (1985), 'Finite element calculations of very high Rayleigh number thermal convection', *Geophys. J. Roy. Astr. Soc.* **80**, 575–601.
- Schubert, G., Anderson, C. & Goldman, P. (1995), 'Mantle plume interaction with an endothermic phase change', *J. Geophys. Res.* **100**, 8245–8256.
- Schubert, G., Turcotte, D. L. & Olson, P. (2001), *Mantle Convection in the Earth and Planets*, Cambridge University Press, Cambridge.
- Scott, D. R. (1992), 'Small-scale convection and mantle melting beneath mid-ocean ridges, in Mantle Flow and Melt Generation at Mid-Ocean Ridges', *Geophys. Monogr. Ser.* **71**, 327–352.
- Sidorin, I., Gurnis, M. & Helmberger, D. V. (1999), 'Dynamics of a phase change at the base of the mantle consistent with seismological observations', *J. Geophys. Res.* **104**, 15005–15023.
- Silver, P. G. & Behn, M. D. (2008), 'Intermittent plate tectonics?', *Science* **319**, 85–88.
- Stacey, F. (1992), *Physics of the Earth*, Brookfield, Brisbane, Queensland, Australia.
- Stefanick, M. & Jurdy, D. M. (1984), 'The distribution of hotspots', *J. Geophys. Res.* **89**, 9919–9925.
- Stegman, D. R., Jellinek, A. M., Zatman, S. A., Baumgardner, J. R. & Richards, M. A. (2003), 'An early lunar core dynamo driven by thermochemical mantle convection', *Nature* **421**, 143–146.
- Stein, C. A. & Stein, S. (1994), 'Constraints on hydrothermal heat-flux through the oceanic lithosphere from global heat flow', *J. Geophys. Res.* **99**, 3081–3095.
- Steinberger, B. & O'Connell, R. J. (1998), 'Advection of plumes in mantle flow: Implications for hotspot motion, mantle viscosity and plume distribution', *Geophys. J. Int.* **132**, 412–434.

- Steinberger, B., Sutherland, R. & O'Connell, R. J. (2004), 'Prediction of Emperor-Hawaii seamount locations from a revised model of global plate motion and mantle flow', *Nature* **430**, 167–173.
- Stemmer, K., Harder, H. & Hansen, U. (2006), 'A new method to simulate convection with strongly temperature- and pressure-dependent viscosity in a spherical shell: Applications to the Earth's mantle', *Phys. Earth Planet. Int.* **157**, 223–249.
- Stixrude, L., Wasserman, E. & Cohen, R. E. (1997), 'Composition and temperature of Earth's inner core', *J. Geophys. Res.* **102**, 24729–24739.
- Tackley, P. (2000a), 'Self-consistent generation of tectonic plates in time-dependent, three-dimensional mantle convection simulations; 1. Pseudoplastic yielding', *Geochem. Geophys. Geosys.* **1**.
- Tackley, P. (2000b), 'Self-consistent generation of tectonic plates in time-dependent, three-dimensional mantle convection simulations; 2. Strain weakening and asthenosphere', *Geochem. Geophys. Geosys.* **1**.
- Tackley, P. (2006), 'Heating up the hotspot debates', *Science* **313**, 1240–1241.
- Tackley, P. J. (1996), 'Effects of strongly variable viscosity on three-dimensional compressible convection in planetary mantles', *J. Geophys. Res.* **101**, 3311–3332.
- Tackley, P. J. (1998), 'Three-dimensional simulations of mantle convection with a thermochemical CMB boundary layer: D'? in: *The Core-Mantle Boundary Region*, *Geodyn. Ser.* **28**, 231–253.
- Tackley, P. J. (2000c), 'Mantle convection and plate tectonics: Toward an integrated physical and chemical theory', *Science* **288**, 2002–2007.
- Tackley, P. J. & King, S. D. (2003), 'Testing the tracer ratio method for modeling active compositional fields in mantle convection simulations', *Geochem. Geophys. Geosys.* **4**, 8302, doi:10.1029/2001GC000214.
- Tackley, P. J., Stevenson, D. J., Glatzmaier, G. A. & Schubert, G. (1994), 'Effects of multiple phase-transitions in a 3-D spherical model of convection in Earth's mantle', *J. Geophys. Res.* **99**, 15877–15901.
- Tan, E., Gurnis, M. & Han, L. J. (2002), 'Slabs in the lower mantle and their modulation of plume formation', *Geochem. Geophys. Geosys.* **3**, 1067, doi:10.1029/2001GC000238.
- Tapley, B., Ries, J., Bettadpur, S., Chambers, D., Cheng, M., Condi, F., Gunter, B., Kang, Z., Nagel, P., Pastor, R., Pekker, T., Poole, S. & Wang, F. (2005), 'GGM02 - An improved Earth gravity field model from GRACE', *J. Geodesy* **79**, 467–478.
- Tarduno, J. A., Duncan, R. A., Scholl, D. W., Cottrell, R. D., Steinberger, B., Thordarson, T., Kerr, B. C., Neal, C. R., Frey, F. A., Torii, M. & Carvallo, C. (2003), 'The Emperor Seamounts: Southward motion of the Hawaiian hotspot plume in Earth's mantle', *Science* **301**, 1064–1069.

- Tarduno, J. A. & Gee, J. (1995), 'Large-scale motion between Pacific and Atlantic hotspots', *Nature* **378**, 477–480.
- Temam, R. (1977), *Navier-Stokes Equations*, North-Holland, Amsterdam.
- Thomas, C., Hessom, T. & Kendall, J. M. (2002), 'Investigating the heterogeneity of the D" region beneath the northern Pacific using a seismic array', *J. Geophys. Res.* **107**, 2274–2290.
- Thompson, C. P., Leaf, G. K. & Van Rosendale, J. (1992), 'A dynamically adaptive multi-grid algorithm for the incompressible Navier-Stokes equation - validation and model problems', *Appl. Num. Math.* **9**, 511–532.
- Thompson, P. F. & Tackley, P. J. (1998), 'Generation of mega-plumes from the core-mantle boundary in a compressible mantle with temperature-dependent viscosity', *Geophys. Res. Lett.* **25**, 1999–2002.
- Tsuchiya, T., Tsuchiya, J., Umemoto, K. & Wentzcovitch, R. A. (2004), 'Phase transition in MgSiO₃ perovskite in the Earth's lower mantle', *Earth Planet. Sci. Lett.* **224**, 241–248.
- Turcotte, D. L. & Oxburgh, E. R. (1967), 'Finite amplitude convective cells and continental drift', *J. Fluid. Mech.* **28**, 29–42.
- Turcotte, D. L. & Schubert, G. (1982), *Application of continuum physics to geological problems*, Wiley, New York.
- Turcotte, D. L. & Schubert, G. (2002), *Geodynamics*, 2nd edn, Cambridge University Press.
- van der Hilst, R. D., Widiyantoro, S. & Engdahl, E. R. (1997), 'Evidence for deep mantle circulation from global tomography', *Nature* **386**, 578–584.
- van Keken, P. E. (1997), 'Evolution of starting mantle plumes: A comparison between numerical and laboratory models', *Earth Planet. Sci. Lett.* **148**, 1–11.
- van Keken, P. E., Hauri, E. H. & Ballentine, C. J. (2002), 'Mantle mixing: The generation, preservation, and destruction of chemical heterogeneity', *Ann. Rev. Earth Planet. Sci.* **30**, 493–525.
- van Keken, P. E., King, S. D., Schmeling, H., Christensen, U. R., Neumeister, D. & Doin, M. P. (1997), 'A comparison of methods for the modeling of thermochemical convection', *J. Geophys. Res.* **102**, 22477–22495.
- Verfuerth, R. (1984), 'A combined Conjugate Gradient-multigrid algorithm for the numerical solution of the Stokes problem', *IMA J. Num. Anal.* **4**, 441–455.
- Wagner, C. (1994), 'A multigrid method with matrix-dependent transfer operators applied in groundwater science', *Comp. Meth. Wat. Resourc.* **2**, 1439–1446.
- Weatherill, N. P. (1989), 'Mesh generation in computational fluid dynamics, *von Karman Institute for Fluid Dynamics*, Lecture Series 1989-04, Brussels'.

- Weidenschilling, S. J., Spaute, D., Davis, D. R., Marzari, F. & Ohtsuki, K. (1997), 'Accretional evolution of a planetesimal swarm .2. The terrestrial zone', *ICARUS* **128**, 429–455.
- Weinstein, S. A. & Olson, P. (1990), 'Planforms in thermal convection with internal heat sources at large Rayleigh and Prandtl numbers', *Geophys. Res. Lett.* **18**, 239–242.
- Weinstein, S. A. & Olson, P. L. (1992), 'Thermal convection with non-newtonian plates', *Geophys. J. Int.* **111**, 515–530.
- Wessel, P., Harada, Y. & Kroenke, L. W. (2006), 'Toward a self-consistent, high-resolution absolute plate motion model for the Pacific', *Geochem. Geophys. Geosys.* **7**, Q03L12, doi:10.1029/2005GC001000.
- Wesseling, P. (1992), *An introduction to multigrid methods*, Wiley, Chichester.
- Wetherill, G. W. & Stewart, G. R. (1989), 'Accumulation of a swarm of small planetesimals', *ICARUS* **77**, 330–357.
- Whittaker, J. M., Muller, R. D., Leitchenkov, G., Stagg, H., Sdrolias, M., Gaina, C. & Goncharov, A. (2007), 'Major Australian-Antarctic plate reorganization at Hawaiian-Emperor bend time', *Science* **318**, 83–86.
- Wilde, S. A., Valley, J. W., Peck, W. H. & Graham, C. M. (2001), 'Evidence from detrital zircons for the existence of continental crust and oceans on the Earth 4.4 Gyr ago', *Nature* **409**, 175–178.
- Yang, W. S. (1997), Variable viscosity thermal convection at infinite Prandtl number in a thick spherical shell, PhD thesis, University of Illinois, Urbana-Champaign.
- Yang, W. S. & Baumgardner, J. R. (2000), 'Matrix-dependent transfer multigrid method for strongly variable viscosity infinite Prandtl number thermal convection', *Geophys. Astrophys. Fluid Dyn.* **92**, 151–195.
- Yoshida, M. & Kageyama, A. (2004), 'Application of the Yin-Yang grid to a thermal convection of a Boussinesq fluid with infinite Prandtl number in a three-dimensional spherical shell', *Geophys. Res. Lett.* **31**.
- Zhang, S. X. & Yuen, D. A. (1995), 'The influences of lower mantle viscosity stratification on 3-D spherical-shell mantle convection', *Earth Planet. Sci. Lett.* **132**, 157–166.
- Zhong, S. J. (2006), 'Constraints on thermochemical convection of the mantle from plume heat flux, plume excess temperature, and upper mantle temperature', *J. Geophys. Res.* **111**, B04409.
- Zhong, S. J. & Gurnis, M. (1995), 'Mantle convection with plates and mobile, faulted plate margins', *Science* **267**, 838–843.
- Zhong, S. J., Zuber, M. T., Moresi, L. & Gurnis, M. (2000), 'Role of temperature-dependent viscosity and surface plates in spherical shell models of mantle convection', *J. Geophys. Res.* **105**, 11063–11082.

-
- Zienkiewicz, O. C. & Taylor, R. L. (2000), *The Finite Element Method*, Vol. 1, 4th edn, McGraw-Hill, Maidenhead.
- Zienkiewicz, O. C. & Zhu, J. Z. (1987), 'A simple error estimator and adaptive procedure for practical engineering analysis', *Int. J. Num. Meth. Eng.* **24**, 337–357.

Appendix A

The Finite Element Method

A.1 Introduction

The process of subdividing or *discretizing* a system into components or ‘elements’ whose behavior is understood and then reconstructing the system from such components to investigate its behavior is a natural way in which the engineer, the scientist or even the economy proceeds. An adequate understanding of a system can often be obtained through models based upon a finite number of components. Such models are termed *discrete* and the variables are clearly distinct from each other.

The advent of powerful digital computers means that *discrete* problems can now be solved with relative ease. *Continuum* problems on the other hand, where one value of the variable flows into the next, are far less simple, yet these are the problems that engineers and scientists must solve everyday. Examples of such situations can be found in the cooling of electronic equipment, the dispersion of pollutants during non-uniform atmospheric conditions and large scale problems such as global climate change. Various methods of discretization have been proposed to help increase our understanding of these systems. All

involve an approximation to the exact solution, which hopefully approaches in the limit, i.e. as the number of discrete variables increases, the true continuum solution. Examples include finite difference approximations (McKenzie et al. 1974, Bodri & Bodri 1978), a well-established and conceptually simple method that requires a point-wise approximation to the governing equations on a meshed geometry, and the finite volume method (Tackley 1998, Ratcliff et al. 1998), a revised version of the finite difference method, based upon small control volumes surrounding each node on a mesh, and inter-volume fluxes. The finite element method (FEM) is an alternative numerical procedure that was originally introduced by engineers in the 1950's for the solution of problems in structural mechanics. However, it soon became apparent that the approach formed the basis of a general technique for the numerical solution of partial differential equations and it has since been widely applied in many branches of computational mechanics, including computational geodynamics (Baumgardner 1985, Farnetani & Richards 1995, Zhong et al. 2000). The methods diversity and flexibility mean that it is now viewed as a general discretization procedure of continuum problems posed by mathematically defined statements.

A.2 Overview

The FEM considers that the solution region comprises many small, interconnected, sub-regions or *elements* and gives a piece-wise approximation to the governing equations. That is, complex partial differential equations are reduced to either linear or non-linear simultaneous equations. Thus the finite element discretization (i.e. dividing the region into a number of smaller regions) procedure reduces the continuum problem, which has an infinite number of unknowns, to one with a finite number of unknowns at specified points or *nodes*. Since the FEM allows us to form the elements, or sub-regions, in an arbitrary sense, a close




Figure A.1: A flow diagram summarizing the major components involved in the finite element solution algorithm.

representation of the boundaries of complicated domains is possible.

This appendix provides an overview of the fundamental components of the FEM. It will be useful not only for readers of this thesis, but also for engineers and scientists seeking to develop an understanding of the method.

A numerical model starts with a physical model of the problem (see Figure A.1). As can be seen, one section of the finite element solution algorithm deals with the discrete approximation of the partial differential equations (equation discretization), while the other deals with domain discretization (spatial discretization). Finally by combining both, the numerical solution to the problem is achieved.

A.3 Equation Discretization

A.3.1 Mathematical Preliminaries

The FEM is based around the discrete representation of a weak integral form of the partial differential equation to be solved. The formulation and subsequent discretization of such an integral form requires the definition of some function spaces and associated norms.

Consider a spatial region (or domain), $\Omega \subset \mathbb{R}^{n_{sd}}$, with piecewise smooth boundary Γ . Here, $n_{sd} = 1, 2$ or 3 , denotes the number of spatial dimensions. We shall use the notation:

$$f : \bar{\Omega} \rightarrow \mathbb{R} \tag{A.1}$$

to state that for each spatial point $\mathbf{x} \in \bar{\Omega}$, $f(\mathbf{x}) \in \mathbb{R}$. $\bar{\Omega}$ denotes the closure of Ω , that is the union of the domain Ω with its boundary Γ : $\bar{\Omega} = \Omega \cup \Gamma$.

A function $f : \bar{\Omega} \rightarrow \mathbb{R}$ is said to be of class $C^m(\Omega)$ if all its derivatives up to order m exist and are continuous functions. For instance, the notation $f(x) \in C^m([a, b])$ indicates that $f(x)$ possesses m continuous derivatives for $x \in]a, b[$.

In finite element analysis, one works with integral equations and, thus, we are interested in functions belonging to larger spaces than C^m . As will be seen, instead of requiring the m -th derivative be continuous, we will require that its square is integrable. In fact, finite element functions should possess generalized derivatives (i.e. derivatives in the sense of distributions) and some integrability properties. Such classes of functions are particular examples of Sobolev function spaces.

Some Useful Sobolev Spaces

Let $L_2(\Omega)$ denote the space of square-integrable functions defined on Ω . This space is equipped with the standard inner product:

$$(u, v) = \int_{\Omega} uv \, d\Omega \quad (\text{A.2})$$

and norm:

$$\|v\|_0 = (v, v)^{1/2} \quad (\text{A.3})$$

For any non-negative integer k , we define the Sobolev space $H^k(\Omega)$ using multi-index notation: given the n -tuple $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_{n_{sd}}) \in \mathbb{N}^{n_{sd}}$ and the non-negative integer $|\alpha| := \alpha_1 + \alpha_2 + \dots + \alpha_{n_{sd}}$,

$$H^k(\Omega) = \left\{ u \in L_2(\Omega) \left| \frac{\partial^{|\alpha|} u}{\partial x_1^{\alpha_1} \partial x_2^{\alpha_2} \dots \partial x_{n_{sd}}^{\alpha_{n_{sd}}}} \in L_2(\Omega) \, \forall |\alpha| \leq k \right. \right\} \quad (\text{A.4})$$

Therefore, $H^k(\Omega)$ consists of square integrable functions all of whose derivatives of order up to k are also square integrable. $H^k(\Omega)$ is equipped with the norm:

$$\|u\|_k = \left(\sum_{s=0}^k \sum_{|\alpha|=s} \left\| \frac{\partial^{|\alpha|} u}{\partial x_1^{\alpha_1} \partial x_2^{\alpha_2} \dots \partial x_{n_{sd}}^{\alpha_{n_{sd}}}} \right\|_0^2 \right)^{1/2} \quad (\text{A.5})$$

Note that L_2 is a Sobolev space, $H^0(\Omega) = L_2(\Omega)$, while the Sobolev space for $k = 1$ is defined by:

$$H^1(\Omega) = \left\{ v \in L_2(\Omega) \left| \frac{\partial v}{\partial x_i} \in L_2(\Omega) \, i = 1, \dots, n_{sd} \right. \right\} \quad (\text{A.6})$$

This space is equipped with the inner product:

$$(u, v)_1 = \int_{\Omega} \left(uv + \sum_{i=1}^{n_{sd}} \frac{\partial u}{\partial x_i} \frac{\partial v}{\partial x_i} \right) d\Omega \quad (\text{A.7})$$

and its induced norm:

$$\|u\|_1 = \sqrt{(u, u)_1} \quad (\text{A.8})$$

The subspace:

$$H_0^1(\Omega) = \{v \in H^1(\Omega) | v = 0 \text{ on } \Gamma\} \quad (\text{A.9})$$

is used frequently in finite element analysis, the elements of which possess a square integrable first derivative over the domain Ω and vanish on its boundary Γ . Moreover, its inner product and norm coincide with those of $H^1(\Omega)$.

It should be noted that the Sobolev spaces used in the remainder of this appendix, namely H^0 , H^1 and H_0^1 , are Hilbert spaces (an abstract vector space in which distances and angles can be measured, which is ‘complete’, meaning that if a sequence of vectors approaches a limit, then that limit is guaranteed to be in the space as well).

Extension to Vector Functions

In the forthcoming sections, consideration will not only be given to scalar functions, but also to vector-valued functions. For vector-valued functions with m components, that is, $\mathbf{u}, \mathbf{v} : \Omega \rightarrow \mathbb{R}^m$, the procedure is in fact essentially the same as for scalar functions.

Consider again a domain $\Omega \subset \mathbb{R}^{n_{sd}} \geq 1$, and denote by $\mathbf{H}^k(\Omega)$ the space of vector functions with m components:

$$\mathbf{u} = (u_1, u_2, \dots, u_m) \quad (\text{A.10})$$

for which each component $u_i \in H^k(\Omega)$, $1 \leq i \leq m$. The space $\mathbf{H}^k(\Omega)$ is equipped with an inner product inducing the following norm:

$$\|\mathbf{u}\|_k = \left(\sum_{i=1}^m \|u_i\|_k^2 \right)^{1/2} \quad (\text{A.11})$$

For the particular case of functions belonging to $\mathbf{L}_2(\Omega) = \mathbf{H}^0(\Omega)$, the inner product is given by:

$$(\mathbf{u}, \mathbf{v}) = \int_{\Omega} \mathbf{u} \cdot \mathbf{v} \, d\Omega \quad (\text{A.12})$$

where there should be no ambiguity in using the same notation to represent the inner product of both scalar and vector-valued functions.

A.3.2 Trial Solutions and Weighting Functions

To define the weak, or variational, form of the boundary value problems (a differential equation, to be satisfied at all points in the interior of an interval, and a set of boundary conditions specifying the values of the solution or some of its derivatives everywhere on the boundary of the interval) discussed in the present text, we need to define two classes of functions: the *test* or weighting functions and the *trial* solutions. Here, these spaces are defined in the context of a standard Galerkin formulation, which will be discussed in more detail later.

The first collection of functions, denoted by V , is composed of *test* functions and consists of all functions which are square integrable, have square integrable first derivatives over the computational domain Ω , and vanish on the Dirichlet portion, Γ_D , of the boundary. They are defined as follows:

$$V = \{w \in H^1(\Omega) \mid w = 0 \text{ on } \Gamma_D\} \equiv H_{\Gamma_D}^1(\Omega) \quad (\text{A.13})$$

As previously noted, this is a Sobolev space and its inner product and norm coincide with those of $H^1(\Omega)$.

The second collection of functions are called the *trial* solutions. This collection is similar to the test functions except that these admissible functions are required to satisfy the Dirichlet conditions on Γ_D . This second collection is denoted by S and is defined by:

$$S = \{u \in H^1(\Omega) \mid u = u_D \text{ on } \Gamma_D\} \equiv V + \{\bar{u}_D\} \quad (\text{A.14})$$

where $\bar{u}_D$ is any function in $H^1(\Omega)$ such that $\bar{u}_D = u_D$ on Γ_D .

The sets S and V contain infinitely many functions. In the FEM, S and V are approximated by convenient finite dimensional subsets of these collections, which will be denoted by S^h and V^h , respectively. These finite element spaces are characterized, amongst other things, by a partition of the domain. The process of domain discretization is described in Section A.5. For now, it is sufficient to assume that the domain has been appropriately partitioned. The reader should view the domain Ω as discretized into elements. Let $T^h(\Omega)$ be a partition of Ω into n_{el} convex subdomains/elements, such that:

$$\bar{\Omega} = \bigcup_{e=1}^{n_{el}} \bar{\Omega}^e \quad \text{and} \quad \bar{\Omega}^e \cap \Omega^f = \emptyset \quad \text{for} \quad e \neq f \quad (\text{A.15})$$

Each subdomain Ω^e has a piecewise smooth boundary Γ_e and h is a characteristic mesh/partition size.

The weighting functions $w^h \in V^h$ vanish on Γ_D . The approximation u^h lies in S^h and satisfies, with precision given by the mesh size h , the boundary condition u_D on Γ_D . In fact, along Γ_D , we should have $u^h = u_D^h$; however, in order not to overload the notation u_D is also used instead of u_D^h unless it is necessary to explicitly show the precision of the boundary data. S^h and V^h should now be thought of as being subsets of S and V , respectively:

$$s^h \subset S \quad (\text{i.e., if } u^h \in S^h, \text{ then } u^h \in S) \quad (\text{A.16})$$

$$v^h \subset V \quad (\text{i.e., if } w^h \in V^h, \text{ then } w^h \in V) \quad (\text{A.17})$$

where the precise meaning is given in parentheses.

A.3.3 Strong and Weak Forms of a Boundary Value Problem

The FEM rests upon an integral form of the partial differential equations under consideration. This was the reason behind our introduction to larger spaces than C^m based on integrability properties. Thus, the first task in a finite element analysis consists of formulating a variational problem associated with the given partial differential equation and its

boundary conditions. At this point a model boundary value problem is introduced. This will be used to illustrate the various steps in the practical implementation of the FEM. Consider solving the Poisson equation:

$$-\nabla^2 u = s \quad \text{in } \Omega \quad (\text{A.18})$$

where Ω is enclosed by a piecewise smooth boundary Γ , and $s \in C^0(\Omega)$ is a specified source term, which may depend on x . We further assume that the value of the unknown u is prescribed on the Dirichlet portion, Γ_D , of the boundary,

$$u = u_D \quad \text{on } \Gamma_D \quad (\text{A.19})$$

while the normal derivative of u is prescribed on the remaining Neumann portion, Γ_N , of the boundary Γ :

$$n \cdot \nabla u = h \quad \text{on } \Gamma_N \quad (\text{A.20})$$

A function $u \in C^2(\Omega) \cap C^0(\Omega)$ that satisfies these equations is called a *classical solution* of the boundary value problem.

The first step in a *weighted residual formulation* leading to the finite element approximation consists of formulating a *weak* (or *variational*) form of the boundary value problem. This is achieved by multiplying the governing equation (A.18) by a suitable weighting function, w , and integrating over the computational domain Ω :

$$-\int_{\Omega} w \nabla^2 u \, d\Omega = \int_{\Omega} w s \, d\Omega \quad (\text{A.21})$$

Note that the continuity requirements on u still impose that it must be twice differentiable. Now, however, the second derivatives of u are not required to be continuous, they only need be square integrable. Thus, it is sufficient that $u \in H^2(\Omega)$. In any case, it is obvious that classical solutions of Equation (A.18) will also verify this integral equation for all admissible functions w .

In order to derive the weak form and produce a natural Neumann boundary condition on Γ_N , we integrate by parts the first term on the left hand side of the previous equation and then apply the Green-Gauss divergence theorem:

$$\begin{aligned} - \int_{\Omega} w \nabla^2 u \, d\Omega &= - \int_{\Omega} (\nabla \cdot (w \nabla u) - \nabla w \cdot \nabla u) \, d\Omega \\ &= \int_{\Omega} \nabla w \cdot \nabla u \, d\Omega - \int_{\Gamma} w (n \cdot \nabla u) \, d\Gamma \end{aligned} \quad (\text{A.22})$$

Now the regularity requirements on w and u are modified; u is only differentiated once and w must be differentiable. In fact, its derivatives must be square integrable; thus we should have $w \in H^1(\Omega)$. If $w \in V$, recall from Equation (A.13) that $w = 0$ on Γ_D . Bearing this in mind and taking into account the prescribed Neumann boundary condition (A.20), we obtain the following weak form of our model problem:

$$\int_{\Omega} \nabla w \cdot \nabla u \, d\Omega = \int_{\Omega} w s \, d\Omega + \int_{\Gamma_N} w h \, d\Gamma \quad (\text{A.23})$$

Note that the application of the divergence theorem has allowed us to naturally introduce the Neumann boundary condition on Γ_N . At the same time, it has removed the second-derivative terms of the Laplacian operator from the volume integral. This reduces the continuity requirements on u and allows us to select this function in $H^1(\Omega)$, since only first derivatives appear under the integral sign.

Classical solutions of Equation (A.18) will also satisfy this integral equation, however, such a solution would only account for Equations (A.18) and (A.20). The Dirichlet boundary conditions (A.19) would not be satisfied. The weak form (A.23) overcomes this through the correct choice of space to which u belongs, namely S . Recall from Equation (A.14) that every member of S satisfies the Dirichlet boundary condition $u = u_D$ on Γ_D . Provided $u \in S$, the Dirichlet boundary condition is satisfied. The weak form of our model problem

can then be formally stated as:

$$\text{Find } u \in S \text{ such that } a(w, u) = (w, s) + (w, h)_{\Gamma_N} \quad \forall w \in V. \quad (\text{A.24})$$

Here, use has been made of bilinear forms:

$$a(w, u) = \int_{\Omega} \nabla w \cdot \nabla u \, d\Omega \quad (\text{A.25})$$

as well as their linear functionals:

$$(w, s) = \int_{\Omega} w s \, d\Omega, \quad \text{and} \quad (w, h)_{\Gamma_N} = \int_{\Gamma_N} w h \, d\Gamma \quad (\text{A.26})$$

If we define u_0 , the solution away from the Dirichlet portion of the domain boundary, Γ_D , as:

$$u = u_0 + \bar{u}_D \quad \text{with} \quad u_0 \in V \quad (\text{A.27})$$

The final weak form of the governing differential equations becomes:

$$\text{Find } u_0 \in V \text{ such that } a(w, u_0) = (w, s) + (w, h)_{\Gamma_N} - a(w, \bar{u}_D) \quad \forall w \in V. \quad (\text{A.28})$$

The right-hand side consists of the totality of terms associated with given data.

A.3.4 Galerkin's Approximation Method

We now have all the necessary ingredients to discretize the weak form by means of the Galerkin FEM. For a domain partition T^h (see A.3.2 and A.5), the Galerkin formulation of our model problem is obtained by restricting the weak form to the finite-dimensional spaces $S^h \subset S$ and $V^h \subset V$ defined in Equations (A.16) and (A.17):

$$\text{Find } u_0^h \in V^h \text{ such that } a(w^h, u_0^h) = (w^h, s) + (w^h, h)_{\Gamma_N} - a(w, \bar{u}_D^h) \quad \forall w^h \in V^h. \quad (\text{A.29})$$

Computational Aspects

Due to the presence of the Dirichlet boundary conditions, a distinction must be made between the number of nodal points, n_{np} , of the discretized domain and the number of nodal unknowns, that is, the number of equations, n_{eq} , of the system resulting from the Galerkin approximation. Following the terminology introduced by Hughes (2000), we denote by $\eta = \{1, 2, \dots, n_{np}\}$ the set of global node numbers in the finite element mesh. Furthermore, we denote by $\eta_D \subset \eta$ the subset of nodes belonging to the Dirichlet portion of the boundary. It follows, in scalar problems, that the cardinal of $\eta \setminus \eta_D$ is equal to n_{eq} , the number of equations.

With this notation, the approximation u^h can be written as:

$$u^h(\mathbf{x}) = \sum_{A \in \eta \setminus \eta_D} N_A(\mathbf{x}) u_A + \sum_{A \in \eta_D} N_A(\mathbf{x}) u_D(\mathbf{x}_A) \quad (\text{A.30})$$

where N_A is the shape or basis function associated with node number A in the finite element mesh, and u_A is the nodal unknown. Moreover, in the Galerkin formulation, the arbitrary test functions, w^h , are defined such that:

$$w^h \in V^h = N_A \quad \text{for } A \in \eta \setminus \eta_D \quad (\text{A.31})$$

Thus, using the definitions of Equations (A.29), (A.30) and (A.31), we obtain the following discrete weak form of our model problem:

$$\sum_{A \in \eta \setminus \eta_D} a(N_A, N_B) u_B = (N_A, s) + (N_A, h)_{\Gamma_N} - \sum_{B \in \eta_D} a(N_A, N_B) u_D(\mathbf{x}_B), \quad \forall A \in \eta \setminus \eta_D \quad (\text{A.32})$$

It should be noted that upper-case letters, such as A and B , are used to represent global node numbers in the finite element mesh: $1 \leq A, B \leq n_{np}$. Lower-case letters, such as a and b , will be used to represent local node numbers in an element: $1 \leq a, b \leq n_{en}$, where n_{en} is the number of element nodes.




Figure A.2: Representative unstructured finite element quadrilateral mesh in 2-D.

A suitable mesh generator is used to subdivide the computational domain Ω into element domains, Ω^e . This process is further described in section A.5. An example of an unstructured, two-dimensional quadrilateral mesh is shown in Figure A.2.

In the practical implementation of the finite element method, attention is focussed on the computations in an individual element. For every element $\Omega^e \in T^h$, the shape functions N_a , $a = 1, \dots, n_{en}$, are defined on a master element in terms of normalized coordinates. They define the variation on a variable over an element and are usually represented by polynomial functions. Let us consider a subdivision of the computational domain Ω into four-node quadrilaterals (see Figure A.3). Each quadrilateral is mapped onto a canonical square with normalized local coordinates, $(\xi, \eta) \in [-1, 1] \times [-1, 1]$, and the element shape functions are tensor products of those used in one dimension. This leads to a so-called iso-parametric bilinear approximation on each element where both the global coordinates (x, y) and the unknown u are expressed by the same bilinear expansion parameterized by




Figure A.3: Four-node quadrilateral elements and normalized reference element.

the nodal values:

$$\begin{Bmatrix} x \\ y \end{Bmatrix} = \sum_{a=1}^4 N_a(\xi, \eta) \begin{Bmatrix} x_a \\ y_a \end{Bmatrix}, \quad \text{and} \quad u^h(x, y) \equiv u^h(\xi, \eta) = \sum_{a=1}^4 N_a(\xi, \eta) u_a \quad (\text{A.33})$$

where (x_a, y_a) are the coordinates of node a of the element. All element contributions to the discrete weak form (A.32) are computed in the local coordinates (ξ, η) using a numerical integration procedure such as Gauss quadrature (see Hughes 2000, Zienkiewicz & Taylor 2000, for details).

The assembly of the element contributions to the discrete weak form into the complete system results in a matrix equation of the form:

$$\mathbf{K} \mathbf{u} = \mathbf{f} \quad (\text{A.34})$$

where $\mathbf{u}$ is the vector of unknown nodal values, of dimension n_{eq} . In practice, the global matrix $\mathbf{K}$ and nodal vector $\mathbf{f}$, result from the topological assembly of element contributions. The addition of the element contributions to the appropriate locations can be represented

through the action of an assembly operator, $\mathbf{A}^e$ acting on the local element matrix and nodal vector as follows:

$$\mathbf{K} = \mathbf{A}^e \mathbf{K}^e, \quad K_{ab}^e = \int_{\Omega^e} \nabla N_a \cdot \nabla N_b d\Omega = a(N_a, N_b)_{\Omega^e} \quad (\text{A.35})$$

$$\mathbf{f} = \mathbf{A}^e \mathbf{f}^e, \quad f_a^e = (N_a, s)_{\Omega^e} + (N_a, h)_{\partial\Omega^e \cap \Gamma_N} - \sum_{b=1}^{n_{en}} a(N_a, N_b)_{\Omega^e} u_{Db}^e \quad (\text{A.36})$$

Here, $u_{Db}^e = u_D(\mathbf{x}_b^e)$ if u_D is prescribed at node number b and equals zero otherwise. A detailed explanation on the topological assembly of the matrices and nodal vectors arising from the Galerkin finite element discretization can be found in Hughes (2000).

A.4 Elements and Shape/Interpolation Functions

The finite element shape functions introduced in the previous section will now be described in further detail. By dividing the solution region into a number of elements and approximating the solution over these regions by a suitable known function, a relation between the differential equations and the elements is established. The functions employed to represent the nature of the solution within each element are called *shape functions*, *interpolation functions* or *basis functions*. Polynomial functions are most widely used since they can be integrated or differentiated easily, and the accuracy of the results can be improved by increasing the order of the polynomial (see Figure A.4)

The shape functions for some simple elements are now described in detail. Although only three examples are presented, the general rules are identical for any element in 1-, 2- or 3-D.




Figure A.4: One-dimensional finite elements. (a) A linear element, (b) a quadratic element, (c) linear and (d) quadratic variation of a variable over an element.

A.4.1 One-dimensional Linear Elements

A large number of scientific and engineering problems can be approximated using a one-dimensional finite element model. Examples include channel flow, pipe flow and thermal conduction through a wall. Figure A.4 displays the temperature profile in an element as represented by linear and quadratic polynomials. Let us consider a typical linear element with end nodes ' i ' and ' j ' and corresponding temperatures T_i and T_j . The linear temperature variation in the element is represented by:

$$T(x) = \alpha_1 + \alpha_2 x \quad (\text{A.37})$$

Here, T is the temperature at any location x while the parameters α_1 and α_2 are constants. Since there are two arbitrary constants in the linear representation, only two nodes are required to determine the values of α_1 and α_2 :

$$\begin{aligned} T_i &= \alpha_1 + \alpha_2 x_i \\ T_j &= \alpha_1 + \alpha_2 x_j \end{aligned} \quad (\text{A.38})$$

From the above equations, we get:

$$\begin{aligned} \alpha_1 &= \frac{T_i x_j - T_j x_i}{x_j - x_i} \\ \alpha_2 &= \frac{T_j - T_i}{x_j - x_i} \end{aligned} \quad (\text{A.39})$$

On substituting the values of α_1 and α_2 into Equation (A.37) we get:

$$T = T_i \left[\frac{x_j - x}{x_j - x_i} \right] + T_j \left[\frac{x - x_i}{x_j - x_i} \right] \quad (\text{A.40})$$

or

$$T = N_i T_i + N_j T_j = \begin{bmatrix} N_i & N_j \end{bmatrix} \begin{Bmatrix} T_i \\ T_j \end{Bmatrix} \quad (\text{A.41})$$

where N_i and N_j are the *shape functions*,

$$\begin{aligned} N_i &= \left[\frac{x_j - x}{x_j - x_i} \right] \\ N_j &= \left[\frac{x - x_i}{x_j - x_i} \right] \end{aligned} \quad (\text{A.42})$$

Equation (A.41) can be rewritten as:

$$T = [\mathbf{N}] \{\mathbf{T}\} \quad (\text{A.43})$$

where

$$[\mathbf{N}] = \begin{bmatrix} N_i & N_j \end{bmatrix} \quad (\text{A.44})$$

Item	Node, i	Node, j	Arbitrary x
N_i	1	0	between 0 and 1
N_j	0	1	between 0 and 1
$N_i + N_j$	1	1	1

Table A.1: Properties of linear shape functions

is the shape function matrix, and:

$$\{\mathbf{T}\} = \begin{Bmatrix} T_i \\ T_j \end{Bmatrix} \quad (\text{A.45})$$

is the vector of unknown temperatures.

From equation (A.41) it is clear that the temperature T at any location x can be calculated using the shape functions N_i and N_j evaluated at x . The shape functions at different locations within an element are shown in Table A.4.1. The shape function assumes a value of unity at the designated node and zero at all other nodes. We also see that the sum of all the shape functions in an element is equal to unity anywhere within the element, including the boundaries. These are two essential requirements of the properties of the shape functions of any element in 1-, 2-, or 3-D. Figure A.5 shows the variation of the shape functions and their derivatives within a linear element. A typical linear variation of temperature is also shown in this figure. As seen, the derivatives of the shape functions are constant within an element.

The temperature gradient can be calculated from Equation (A.41) as:

$$\frac{dT}{dx} = \frac{dN_i}{dx}T_i + \frac{dN_j}{dx}T_j = -\frac{1}{x_j - x_i}T_i + \frac{1}{x_j - x_i}T_j \quad (\text{A.46})$$

or:

$$\frac{dT}{dx} = \begin{bmatrix} -\frac{1}{l} & \frac{1}{l} \end{bmatrix} \begin{Bmatrix} T_i \\ T_j \end{Bmatrix} \quad (\text{A.47})$$




Figure A.5: Variation of shape functions, temperature and derivatives within a linear element

where l is the length of an element (i.e. $x_j - x_i$). Thus we observe that the temperature gradient is constant within an element as the temperature variation is linear.

A.4.2 One-dimensional Quadratic Element

Figure A.4 illustrates that the temperature profile over an element is better approximated using parabolic arcs as opposed to linear segments. The function $T(x)$ would therefore be quadratic in x within each element, taking the form

$$T(x) = \alpha_1 + \alpha_2 x + \alpha_3 x^2 \quad (\text{A.48})$$

There are now three unknowns and, consequently, we need the temperature at one more point, in addition to the two end points of the element. The mid-point is selected, leading to the following equations for the temperature at these three locations:

$$\begin{aligned} T_i &= \alpha_1 \\ T_j &= \alpha_1 + \alpha_2 \frac{l}{2} + \alpha_3 \left(\frac{l}{2}\right)^2 \\ T_k &= \alpha_1 + \alpha_2 l + \alpha_3 l^2 \end{aligned} \quad (\text{A.49})$$

From the above three equations, the following values are obtained for the three constants α_1 , α_2 and α_3 :

$$\begin{aligned} \alpha_1 &= T_i \\ \alpha_2 &= \frac{1}{l}(-3T_i + 4T_j - T_k) \\ \alpha_3 &= \frac{2}{l^2}(T_i - 2T_j + T_k) \end{aligned} \quad (\text{A.50})$$

Substituting α_1 , α_2 and α_3 into Equation (A.48) and collating the coefficients of T_i , T_j and T_k we get:

$$T = T_i \left[1 - \frac{3x}{l} + \frac{2x^2}{l^2}\right] + T_j \left[4\frac{x}{l} - 4\frac{x^2}{l^2}\right] + T_k \left[2\frac{x^2}{l^2} - \frac{x}{l}\right] \quad (\text{A.51})$$

$$T = N_i T_i + N_j T_j + N_k T_k \quad (\text{A.52})$$

Hence, the shape functions for a one-dimensional quadratic element are obtained from Equation (A.51) as follows:

$$\begin{aligned} N_i &= \left[1 - \frac{3x}{l} + \frac{2x^2}{l^2}\right] \\ N_j &= \left[4\frac{x}{l} - 4\frac{x^2}{l^2}\right] \\ N_k &= \left[2\frac{x^2}{l^2} - \frac{x}{l}\right] \end{aligned} \quad (\text{A.53})$$

The variation of temperature and shape functions of a typical quadratic element is shown




Figure A.6: Variation of shape functions and their derivatives over a one-dimensional quadratic element.

in Figure A.6. The first derivative of temperature can now be written as:

$$\frac{dT}{dx} = \frac{dN_i}{dx}T_i + \frac{dN_j}{dx}T_j + \frac{dN_k}{dx}T_k \quad (\text{A.54})$$

or:

$$\frac{dT}{dx} = \left[\frac{4x}{l^2} - \frac{3}{l} \right] T_i + \left[\frac{4}{l} - \frac{8x}{l^2} \right] T_j + \left[\frac{4x}{l^2} - \frac{1}{l} \right] T_k \quad (\text{A.55})$$

From Equation (A.53), $N_i = 1$ at i and 0 at j and k , $N_j = 1$ at j and 0 at i and k and $N_k = 1$ at k and 0 at i and j . It can be easily verified that within an element, the summation over the shape functions is equal to unity, that is:

$$\sum_{i=1}^3 N_i = 1 \quad (\text{A.56})$$

For example, at the point $x = l/4$, the shape function values are:

$$N_i = 1 - \frac{3}{4} + \frac{2}{16} = \frac{6}{16}$$




Figure A.7: A one-dimensional linear element represented by local coordinates.

$$\begin{aligned} N_j &= 1 - \frac{4}{16} = \frac{12}{16} \\ N_k &= \frac{2}{16} - \frac{1}{4} = -\frac{2}{16} \end{aligned} \quad (\text{A.57})$$

It can also be observed that even though the derivatives of the quadratic element are functions of the independent variable x , they will not be continuous at the inter-element nodes. The formulation used here only guarantees the continuity of the function across element boundaries - known as C^0 continuity, where, as noted previously, the subscript indicates that only derivatives of zero order are continuous, that is, only the function is continuous. Elements that also assure the continuity of derivatives across inter-element boundaries, in addition to the continuity of functions, are known as C^1 elements.

C^0 shape functions can be determined in a general way using the Lagrangian polynomial formulae. The one-dimensional $(n - 1)$ th order Lagrange interpolation polynomial is the ratio of two products. For an element with n nodes, $(n - 1)$ order polynomial, the interpolation function is:

$$N_k^e(x) = \prod_{i=1}^n \frac{x - x_i}{x_k - x_i} \quad (\text{A.58})$$

Note that in the above equation, $k \neq i$. For a one-dimensional linear element, the shape functions can be written using Equation (A.58), as $(n = 2)$:

$$N_1 = \frac{x - x_2}{x_1 - x_2} \quad \text{and} \quad N_2 = \frac{x - x_1}{x_2 - x_1} \quad (\text{A.59})$$

Note that N_1 and N_2 are the shape functions corresponding to the two nodes of a one-dimensional linear element (N_i and N_j). If we use local coordinates, as illustrated by Figure A.7 with $x_1 = 0$ and $x_2 = 1$, then the shape functions become:

$$N_i = \left(1 - \frac{x}{l}\right) = L_i \quad \text{and} \quad N_j = \left(\frac{x}{l}\right) = L_j \quad (\text{A.60})$$

where L_i and L_j are the shape functions defined by the local coordinate system. For a one-dimensional quadratic element, the shape functions using Lagrangian multipliers are given as follows:

$$\begin{aligned} N_1 &= \frac{x - x_2}{x_1 - x_2} \frac{x - x_3}{x_1 - x_3} \\ N_2 &= \frac{x - x_1}{x_2 - x_1} \frac{x - x_3}{x_2 - x_3} \\ N_3 &= \frac{x - x_1}{x_3 - x_1} \frac{x - x_2}{x_3 - x_2} \end{aligned} \quad (\text{A.61})$$

If we substitute $x_1 = 0$, $x_2 = l/2$ and $x_3 = l$, in the above equation, we can immediately verify that the resulting equations are identical to the one derived from Equation (A.53). Similarly, cubic elements, or any other one-dimensional higher-order element shape functions, can easily be derived using the Lagrangian interpolation formula.

The important point to take from this discussion is that for the case of quadratic and cubic elements, a better approximation of curved shapes is possible as we have more than two points placed along the boundaries of an element. However, the flip side is a greater computational expense.

A.4.3 Two-dimensional Linear Triangular Elements

When one-dimensional approximations are insufficient, multi-dimensional solution procedures must be employed. The simplest geometric shape that can be employed to approximate irregular surfaces in 2-D is the triangle. It is one of the most popular elements currently used in finite element calculations.




Figure A.8: A linear triangular element.

The two-dimensional linear triangular element is represented by:

$$T(x, y) = \alpha_1 + \alpha_2 x + \alpha_3 y \quad (\text{A.62})$$

where the polynomial is linear in x and y and contains three coefficients. Since a linear triangle has three nodes (see Figure A.8), the values of α_1 , α_2 and α_3 are determined from:

$$\begin{aligned} T_i &= \alpha_1 + \alpha_2 x_i + \alpha_3 y_i \\ T_j &= \alpha_1 + \alpha_2 x_j + \alpha_3 y_j \\ T_k &= \alpha_1 + \alpha_2 x_k + \alpha_3 y_k \end{aligned} \quad (\text{A.63})$$

which results in the following:

$$\begin{aligned} \alpha_1 &= \frac{1}{2A} [(x_j y_k - x_k y_j) T_i + (x_k y_i - x_i y_k) T_j + (x_i y_j - x_j y_i) T_k] \\ \alpha_2 &= \frac{1}{2A} [(y_j - y_k) T_i + (y_k - y_i) T_j + (y_i - y_j) T_k] \\ \alpha_3 &= \frac{1}{2A} [(x_k - x_j) T_i + (x_i - x_k) T_j + (x_j - x_i) T_k] \end{aligned} \quad (\text{A.64})$$

where A is the area of the triangle, given by:

$$2A = \det \begin{pmatrix} 1 & x_i & y_i \\ 1 & x_j & y_j \\ 1 & x_k & y_k \end{pmatrix} = (x_i y_j - x_j y_i) + (x_k y_i - x_i y_k) + (x_j y_k - x_k y_j) \quad (\text{A.65})$$

Substituting the values of α_1 , α_2 and α_3 into Equation (A.63) and collating the coefficients of T_i , T_j and T_k , we get:

$$T = N_i T_i + N_j T_j + N_k T_k = \begin{bmatrix} N_i & N_j & N_k \end{bmatrix} \begin{Bmatrix} T_i \\ T_j \\ T_k \end{Bmatrix} \quad (\text{A.66})$$

where:

$$\begin{aligned} N_i &= \frac{1}{2A}(a_i + b_i x + c_i y) \\ N_j &= \frac{1}{2A}(a_j + b_j x + c_j y) \\ N_k &= \frac{1}{2A}(a_k + b_k x + c_k y) \end{aligned} \quad (\text{A.67})$$

and:

$$\begin{aligned} a_i &= x_j y_k - x_k y_j; & b_i &= y_j - y_k; & c_i &= x_k - x_j \\ a_j &= x_k y_i - x_i y_k; & b_j &= y_k - y_i; & c_j &= x_i - x_k \\ a_k &= x_i y_j - x_j y_i; & b_k &= y_i - y_j; & c_k &= x_j - x_i \end{aligned} \quad (\text{A.68})$$

If we evaluate N_i at node i , where the coordinates are (x_i, y_i) , then we obtain:

$$(N_i)_i = \frac{1}{2A}[(x_j y_k - x_k y_j) + (y_j - y_k)x_i + (x_k - x_j)y_i] = \frac{2A}{2A} = 1 \quad (\text{A.69})$$

Similarly, it can be readily verified that $(N_j)_i = (N_k)_i = 0$. Thus we see that the shape functions have a value of unity at the designated vertex and zero at all other vertices.

The gradients of the temperature T are given by:

$$\begin{aligned}\frac{\partial T}{\partial x} &= \frac{\partial N_i}{\partial x} T_i + \frac{\partial N_j}{\partial x} T_j + \frac{\partial N_k}{\partial x} T_k = \frac{b_i}{2A} T_i + \frac{b_j}{2A} T_j + \frac{b_k}{2A} T_k \\ \frac{\partial T}{\partial y} &= \frac{\partial N_i}{\partial y} T_i + \frac{\partial N_j}{\partial y} T_j + \frac{\partial N_k}{\partial y} T_k = \frac{c_i}{2A} T_i + \frac{c_j}{2A} T_j + \frac{c_k}{2A} T_k\end{aligned}\quad (\text{A.70})$$

It should be noted that both $\partial T/\partial x$ and $\partial T/\partial y$ are constants within an element as b_i , b_j , b_k and c_i , c_j , c_k are constants for a given triangle.

A.4.4 Requirements for Interpolation Functions

A finite element solution will converge to the exact solution, as the number of elements increases, provided that the *compatibility* and *completeness* conditions are met. The shape functions employed must therefore satisfy the following requirements:

1. *For Compatibility:* At element interfaces, the field variable T and any of its partial derivatives, up to one order less than the highest-order derivative appearing in the weighted residual form of the governing equations, must be continuous.
2. *For Completeness:* All uniform states of T and its partial derivatives up to the highest-order appearing in the weighted residual form of the governing equations, should have representation in T , when in the limit the element size decreases to zero.

If the field variables are continuous at the element interfaces, then we have C^0 continuity. If, in addition, the first derivatives are continuous, we have C^1 continuity, and if the second derivatives are continuous, then we have C^2 continuity, and so on. Therefore:

1. *For Compatibility:* At element interfaces, there must be C^r continuity.
2. *For Completeness:* Within an element, we must have C^{r+1} continuity.

where r is the order of the highest occurring derivative.




Figure A.9: Structured quadrilateral mesh discretization of a two-dimensional computational domain.

A.5 Spatial Discretization

An integral part of the FEM involves splitting the computational domain into elements. This subdivision is termed *meshing* or *discretization*. It is often quoted as the most important and most time consuming part of Computational Fluid Dynamics (CFD) analysis. The quality of the grid plays a direct role on the quality of the analysis, regardless of the flow solver used. Furthermore, the solver will be more robust and efficient when using a well constructed mesh. In this section, two alternative strategies for discretizing the computational domain are outlined. We will assume, for brevity, that the mesh is to be produced for a two dimensional domain.

In the first approach (e.g. Weatherill 1989), the domain is divided into a structured assembly of quadrilateral cells. The structure in the mesh is apparent from the fact that each interior nodal point is surrounded by exactly the same number of elements, as shown in Figure A.9. Structured grids enjoy a considerable advantage over other grid methods in that they allow the user a high degree of control; the user places nodes interactively and, consequently, has total freedom when positioning the mesh. In addition, quadrilateral (in 2-D) and hexahedral (in 3-D) elements, which are very efficient at filling space, support a high amount of skewness




Figure A.10: Unstructured triangular mesh discretization of a two-dimensional computational domain.

and stretching before solution accuracy begins to degenerate significantly. This allows the user to naturally condense points in regions of high gradients in the flow and expand out to a less dense packing away from these areas. Furthermore, since the user lays out the elements, the grid is most often flow-aligned, thereby yielding greater accuracy within the solver. Structured flow solvers typically require the lowest amount of memory for a given mesh size and execute faster because they are optimized for the structured layout of the grid. Finally, post processing of the results on a structured block grid is typically a much easier task because the logical grid planes make excellent reference points for examining the flow field and plotting results. The major drawback of structured grids is the time and expertise required to lay out an optimal structure for an entire model, which often comes down to past user experience. Additionally, some geometries, eg. shallow cones and wedges, do not lend themselves to structured block topologies. In these areas, the user is forced to stretch or twist the elements to a degree which drastically affects solver accuracy and performance.

The alternative approach is to divide the computational domain into an unstructured assembly of elements as illustrated in Figure A.10. The notable feature of an unstructured mesh is that the number of elements surrounding a typical interior node is not necessarily




Figure A.11: Connectivity array for an unstructured triangular mesh. Note that element numbers are in red, while node numbers are in black.

constant. Quadrilateral elements could again be used in this context as illustrated by Figure A.2, but we will be focussing upon the use of triangular meshes in this section.

The nodes and elements are numbered and, to get the necessary information on the neighbors, we store the numbers of the nodes which belong to each element (see Figure A.11). From the detail of a typical unstructured mesh shown in Figure A.10 it is apparent that there is no concept of directionality within a mesh of this type and that, therefore, solution techniques based upon this concept will not be directly applicable. The methods normally adapted to generate unstructured meshes are based upon either the Delaunay (Baker 1990) or the Advancing Front approach (Peraire et al. 1987). Discretization methods for the equations of fluid flow which are based upon integral procedures, as is the FEM, are natural candidates for use with unstructured meshes.

The principal advantage of the unstructured approach is that it provides a very powerful tool for discretizing domains of complex shape, especially if triangles are used in 2-D and tetrahedra are used in 3-D. In addition, unstructured mesh methods naturally offer the possibility of incorporating adaptivity (Lohner et al. 1986), which is of extreme importance. Furthermore, unstructured grid generation methods are very automated and, therefore, re-

quire little user time or effort. They also enable the solution of very large and detailed problems in a relatively short period of time. Disadvantages which follow from adopting the unstructured grid approach are that the number of alternative solution algorithms is less than those available with structured grids. Additionally, their computational implementation places large demands on both computer memory and CPU (Formaggia et al. 1988). Further, these algorithms are rather sensitive to the quality of the grid being employed and, consequently, great care must be taken in the generation process.

Appendix B

CONMAN

B.1 Introduction

CONMAN is a finite element program that was written by Arthur Raefsky and Scott King at Caltech in 1989. The code solves the two-dimensional equations of incompressible, infinite Prandtl number thermal convection. An overview of the code is provided here. Further details can be found in King et al. (1990), while the underlying numerical methods and formulations are described by Hughes et al. (1979) and Hughes & Brooks (1979).

B.2 Equations and implementation

The equations for incompressible convection (in dimensionless, vector form) are the equations of Momentum:

$$\nabla^2 \mathbf{u} = -\nabla p + RaT\hat{\mathbf{k}} \quad (\text{B.1})$$

Continuity (mass):

$$\nabla \cdot \mathbf{u} = 0 \quad (\text{B.2})$$

Energy:

$$\frac{\partial T}{\partial t} + \mathbf{u} \cdot \nabla T = \nabla^2 T \quad (\text{B.3})$$

where $\mathbf{u}$ is the dimensionless velocity, T is the dimensionless temperature, p is the dimensionless dynamic pressure, $\hat{\mathbf{k}}$ is the unit vector in the direction of gravity and t is the dimensionless time. In this form, all material properties are combined into one dimensionless parameter, the Rayleigh number, given by:

$$Ra = \frac{\beta \rho g \Delta T d^3}{\kappa \mu} \quad (\text{B.4})$$

where g is the acceleration due to gravity, ρ is density, β is the coefficient of thermal expansion, ΔT is the temperature drop across the domain, d is the domain height, κ is the thermal diffusivity and μ is the dynamic viscosity.

Four equations evolve from the conservation equations above, i.e. continuity, momentum in the x and y Cartesian directions, and energy. These take the form:

Continuity:

$$\frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} = 0 \quad (\text{B.5})$$

Momentum:

$$\begin{aligned} \frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2} &= -\frac{\partial p}{\partial x} \\ \frac{\partial^2 v}{\partial x^2} + \frac{\partial^2 v}{\partial y^2} &= -\frac{\partial p}{\partial y} + Ra\theta\hat{k} \end{aligned} \quad (\text{B.6})$$

Energy:

$$\frac{\partial T}{\partial t} + \left(u \frac{\partial T}{\partial x} + v \frac{\partial T}{\partial y} \right) = \left(\frac{\partial^2 T}{\partial x^2} + \frac{\partial^2 T}{\partial y^2} \right) \quad (\text{B.7})$$

In the above equations, u and v are the components of velocity in the x and y directions respectively. The momentum and energy equations form a simple coupled set of differential equations, although coupling is not strong since ρ is constant, other than in the buoyancy

term of the momentum equation (Boussinesq approximation). The incompressibility (continuity) equation is treated as a constraint on the momentum equation, and incompressibility is enforced in the solution of the momentum equation using a penalty formulation. As the temperatures provide the buoyancy (body force) to drive the momentum equation and as there is no time dependence in the momentum equation, the algorithm to solve the system is a simple one: given an initial temperature field, calculate the resulting velocity field. Use the velocities to advect the temperatures for the next time step and solve for a new temperature field.

B.3 The Momentum Equation

B.3.1 Preliminaries

Let Ω be an open set contained in $\mathbb{R}^n$, $n = 2$, with piecewise smooth boundary Γ . The closure of a set is denoted by a superposed bar (e.g. $\overline{\Omega}$ is the closure of Ω). Vector and tensor fields defined on Ω are written in boldface notation. The Cartesian components of vectors and tensors are written in the standard indicial notation. For example, x_i and u_i are the i th components of the position vector $\mathbf{x}$ and velocity vector $\mathbf{u}$ respectively. The summation convention is employed on repeated indices i , j and k only (e.g. $\tau_{ii} = \tau_{11} + \tau_{22} + \dots + \tau_{nn}$), while a comma is used to denote partial differentiation (e.g. $u_{i,j} = \partial u_i / \partial x_j$, the velocity gradients).

Let L_2 denote the space of square-integrable functions defined on Ω , and let H^1 denote the space of L_2 -functions whose partial derivatives are also in L_2 . The L_2 and H^1 norms (a function that assigns a positive length or size to all vectors in a vector space) are defined

by:

$$\|u\|_0 = \left[\int_{\Omega} uu \, d\Omega \right]^{1/2} \quad (\text{B.8})$$

$$\|u\|_1 = \left[\int_{\Omega} (uu + u_{,i}u_{,i}) \, d\Omega \right]^{1/2} \quad (\text{B.9})$$

Let $\mathbf{L}_2$ ($\mathbf{H}^1$, respectively) denote the space of vector fields whose components are in L_2 (H^1 , respectively). The $\mathbf{L}_2$ and $\mathbf{H}^1$ norms are defined by:

$$\|\mathbf{u}\|_0 = \left[\int_{\Omega} u_i u_i \, d\Omega \right]^{1/2} \quad (\text{B.10})$$

$$\|\mathbf{u}\|_1 = \left[\int_{\Omega} (u_i u_i + u_{i,j} u_{i,j}) \, d\Omega \right]^{1/2} \quad (\text{B.11})$$

B.3.2 Prescribed Data and Formal Statement of the Problem

Let Γ_g and Γ_h be subsets of Γ which satisfy the following conditions:

$$\overline{\Gamma_g \cup \Gamma_h} = \Gamma \quad (\text{B.12})$$

$$\Gamma_g \cap \Gamma_h = \emptyset \quad (\text{B.13})$$

Equation (B.12) reads as ‘the set that contains all elements from Γ_g and Γ_h , but no others’. Equation (B.13) translates to ‘all elements that Γ_g and Γ_h have in common’. Since $\emptyset$ denotes an empty set, such elements do not exist. The momentum equation is solved using the penalty method to enforce incompressibility. The formal statement of the boundary value problem is as follows. Given:

$$f : \Omega \rightarrow \mathbb{R}^n \quad (\text{Body force vector}) \quad (\text{B.14})$$

$$g : \Gamma_g \rightarrow \mathbb{R}^n \quad (\text{Imposed velocity vector}) \quad (\text{B.15})$$

$$h : \Gamma_h \rightarrow \mathbb{R}^n \quad (\text{Imposed traction vector}) \quad (\text{B.16})$$

where Γ is the boundary of the domain Ω , Γ_g and Γ_h are the parts of the boundary where velocities (Dirichlet) and tractions (Neumann) are specified.

Find $\mathbf{u} : \Omega \rightarrow \mathbb{R}^n$ and $p : \Omega \rightarrow \mathbb{R}$ such that:

$$\tau_{ij,j} + f_i = 0 \quad \text{on } \Omega \quad (\text{B.17})$$

$$u_{i,i} = 0 \quad \text{on } \Omega \quad (\text{B.18})$$

$$u_i = g_i \quad \text{on } \Gamma_g \quad (\text{B.19})$$

$$\tau_{ij}n_j = h_i \quad \text{on } \Gamma_h \quad (\text{B.20})$$

Since we are considering flows of a Newtonian fluid, the constitutive equation utilized is:

$$\tau_{ij} = -p\delta_{ij} + 2\mu u_{(i,j)} \quad (\text{B.21})$$

where τ_{ij} denotes the Cauchy stress tensor, p is the pressure, δ_{ij} is the Kronecker delta (if $i = j$, then $\delta_{ij} = 1$, otherwise $\delta_{ij} = 0$) and $u_{(i,j)} = (u_{i,j} + u_{j,i})/2$. Note that if $\Gamma_h = \emptyset$ then a consistency condition is needed, emanating from Equations (B.18) and (B.19):

$$0 = \int_{\Omega} u_{i,i} d\Omega \quad (\text{B.22})$$

$$= \int_{\Gamma} u_i n_i d\Gamma \quad \text{Green-Gauss divergence theorem} \quad (\text{B.23})$$

$$= \int_{\Gamma} g_i n_i d\Gamma \quad \text{By Equation (B.19)} \quad (\text{B.24})$$

B.3.3 Penalty Function Formulation

In the penalty-function formulation, the constitutive equation (B.21) is replaced by:

$$\tau_{ij}^{(\lambda)} = -p^{(\lambda)}\delta_{ij} + 2\mu u_{(i,j)}^{(\lambda)} \quad (\text{B.25})$$

in which:

$$p^{(\lambda)} = -\lambda u_{i,i}^{\lambda} \quad (\text{B.26})$$

where λ is the penalty parameter. In this formulation, the continuity and momentum equations are strongly linked. Incompressibility is automatically enforced since the solution

to Equations (B.17), (B.25) and (B.26) converges to the incompressible Stokes flow solution as λ approaches infinity (Temam 1977). An advantage of the method is that the additional unknown p is eliminated from the equation set. This is useful not only because the amount of computational work is decreased since no pressure equation is solved, but also because it eliminates the need to create boundary conditions for the pressure equation. There are therefore no pressure boundary conditions in the formal specification of the problem.

The boundary value problem for the penalty function formulation is stated as follows:

Find $\mathbf{u}^{(\lambda)} : \Omega \rightarrow \mathbb{R}^n$ such that:

$$\tau_{ij,j}^{(\lambda)} + f_i = 0 \quad \text{on } \Omega \quad (\text{B.27})$$

$$u_i^{(\lambda)} = g_i \quad \text{on } \Gamma_g \quad (\text{B.28})$$

$$\tau_{ij}^{(\lambda)} = h_i \quad \text{on } \Gamma_h \quad (\text{B.29})$$

where $\tau_{ij}^{(\lambda)}$ is given by Equation (B.25). The remainder of this section deals solely with the penalty-function formulation. Thus, it is notationally convenient to omit the λ superscripts in all subsequent developments.

B.3.4 Weak Formulation

The first step in the weighted residual formulation leading to the finite element discretization consists of formulating a *weak* or *variational* form of the boundary value problem. This is achieved by multiplying the governing equation by suitable weighting functions and integrating over the computational domain, Ω . Let

$$\mathbf{V} = \{\mathbf{w} \in \mathbf{H}^1 | \mathbf{w} = 0 \quad \text{on } \Gamma_g\} \quad (\text{B.30})$$

Equation (B.30) reads: ‘ $\mathbf{V}$ is the space of all $\mathbf{H}^1$ functions which vanish on Γ_g ’. $\mathbf{V}$ is often called the space of *weighting functions*. Let $\tilde{g}$ denote a given $\mathbf{H}^1$ -extension of g ; that is

$\tilde{g} \in \mathbf{H}^1$ and:

$$\tilde{g} = g \quad \text{on} \quad \Gamma_g \quad (\text{B.31})$$

The weak form of the boundary value problem is stated as follows; find $\mathbf{u} = \mathbf{w} + \tilde{g}$, $\mathbf{w} \in \mathbf{V}$, such that for all $\bar{\mathbf{w}} \in \mathbf{V}$:

$$\begin{aligned} & \int_{\Omega} (\lambda w_{j,j} \bar{w}_{i,i} + 2\mu w_{(i,j)} \bar{w}_{(i,j)}) d\Omega = \\ & \int_{\Omega} \bar{w}_i f_i d\Omega + \int_{\Gamma} \bar{w}_i h_i d\Gamma - \int_{\Omega} (\lambda \tilde{g}_{j,j} \bar{w}_{i,i} + 2\mu \tilde{g}_{(i,j)} \bar{w}_{(i,j)}) d\Omega \end{aligned} \quad (\text{B.32})$$

The first stage in deriving the above relation is to multiply the governing equation by the weighting function $\bar{w}$ and integrate over the computational domain Ω :

$$\int_{\Omega} \bar{w}_i [\lambda u_{i,i} \delta_{ij} + 2\mu u_{(i,j)}]_{,j} d\Omega + \int_{\Omega} \bar{w}_i f_i d\Omega = 0 \quad (\text{B.33})$$

Note that the continuity requirements on $\mathbf{u}$ still impose that it must be twice differentiable. Now, however, the second derivatives of $\mathbf{u}$ are not required to be continuous, they only need to be square integrable. In order to derive the weak form and produce a natural Neumann boundary condition, we integrate by parts the first term on the left hand side then apply the Green-Gauss divergence theorem, which yields:

$$- \int_{\Omega} [\bar{w}_{(i,j)} (\lambda u_{i,i} \delta_{ij} + 2\mu u_{(i,j)})] d\Omega + \int_{\Gamma_h} [\bar{w}_i (\lambda u_{i,i} \delta_{ij} + 2\mu u_{(i,j)})] n_j d\Gamma + \int_{\Omega} \bar{w}_i f_i d\Omega = 0 \quad (\text{B.34})$$

Rearranging:

$$\int_{\Omega} [\bar{w}_{(i,j)} (\lambda u_{i,i} \delta_{ij} + 2\mu u_{(i,j)})] d\Omega = \int_{\Gamma_h} [\bar{w}_i (\lambda u_{i,i} \delta_{ij} + 2\mu u_{(i,j)})] n_j d\Gamma + \int_{\Omega} \bar{w}_i f_i d\Omega \quad (\text{B.35})$$

and utilizing Equation (B.29):

$$\int_{\Omega} [\bar{w}_{(i,j)} (\lambda u_{i,i} \delta_{ij} + 2\mu u_{(i,j)})] d\Omega = \int_{\Gamma} \bar{w}_i h_i d\Gamma + \int_{\Omega} \bar{w}_i f_i d\Omega \quad (\text{B.36})$$

and remembering that $\mathbf{u} = \mathbf{w} + \tilde{g}$, we arrive at Equation (B.32). The application of the divergence theorem has allowed us to naturally introduce the Neumann boundary condition.

At the same time, it has removed the second-derivative terms from the volume integral. This reduces the continuity requirements on $\mathbf{u}$ and allows us to select this function in $\mathbf{H}^1(\Omega)$, since only first derivatives appear under the integral sign.

B.3.5 Galerkin Formulation

The Galerkin counterpart, where weighting functions are the same as the shape functions, of the weak formulation is now presented. $\mathbf{V}^h$ is a subset of $\mathbf{V}$ parameterized by h , the mesh spacing. Let $\tilde{g}^h$ denote an approximation of $\tilde{g}$ which converges to $\tilde{g}$ as $h \rightarrow 0$. Find $\mathbf{u}^h = \mathbf{w}^h + \tilde{g}^h$, $\mathbf{w}^h \in \mathbf{V}^h$, such that for all $\bar{\mathbf{w}}^h \in \mathbf{V}^h$:

$$\begin{aligned} \int_{\Omega} (\lambda w_{j,j}^h \bar{w}_{i,i}^h + 2\mu w_{(i,j)}^h \bar{w}_{(i,j)}^h) d\Omega = \\ \int_{\Omega} \bar{w}_i^h f_i d\Omega + \int_{\Gamma} \bar{w}_i^h h_i d\Gamma - \int_{\Gamma} [(\lambda \tilde{g}_{j,j}^h \bar{w}_{i,i}^h + 2\mu \tilde{g}_{(i,j)}^h \bar{w}_{(i,j)}^h) d\Gamma \end{aligned} \quad (\text{B.37})$$

B.3.6 Matrix Problem

The domain is discretized into quadrilateral elements. The e th element domain is denoted by Ω^e and its boundary is denoted Γ^e . Associated with the discretization is a set of n_{np} nodal points. The position vector of the A th node, $A = 1, 2, \dots, n_{np}$, is denoted $\mathbf{x}_A$. Let $\aleph = \{1, 2, \dots, n_{np}\}$, the set of nodal indices, and let $\aleph_{\tilde{g}} = \{A \in \aleph | \mathbf{x}_A \in \Gamma_{\tilde{g}}\}$, the subset of nodal indices corresponding to boundary nodes at which velocities are prescribed. The ‘shape function’ associated with node A is denoted N_A . The solution of the Galerkin problem may be expressed in terms of the shape functions as follows:

$$w_i^h = \sum_{A \in \aleph - \aleph_{\tilde{g}}} N_A u_{iA} \quad (\text{B.38})$$

$$\tilde{g}_i^h = \sum_{A \in \aleph_{\tilde{g}}} N_A g_{iA} \quad (\text{B.39})$$

where u_{iA} is the i th velocity component at node A and:

$$g_{iA} = g_i(\mathbf{x}_A) \quad (\text{B.40})$$

As can be seen by (B.40), the approximation of $\tilde{g}_i$ assumed in (B.39) is one of nodal interpolation via the shape functions.

Substitution of Equations (B.38) and (B.39) into Equation (B.37) gives rise to the matrix problem:

$$\mathbf{K}\mathbf{u} = \mathbf{f} \quad (\text{B.41})$$

where $\mathbf{K}$ is the element stiffness matrix, $\mathbf{u}$ is the vector of unknown velocities and $\mathbf{f}$ is the right hand side force vector. These are explicitly defined as:

$$\mathbf{K} = [K_{PQ}] \quad (\text{B.42})$$

$$\mathbf{u} = \{u_Q\} \quad (\text{B.43})$$

$$\mathbf{f} = \{f_P\} \quad (\text{B.44})$$

The indices P and Q take on the values 1, 2,..., n_{eq} , where n_{eq} refers to the number of equations in the ‘global’ system. The equation numbers are stored in a ‘destination array’, denoted ID and defined as follows:

$$P = ID(i, A) \quad (\text{B.45})$$

Here, P represents the equation number, i is the degree of freedom number ($1 \leq i \leq n$) and A is the ‘global’ node number. Nodal velocity components, which are prescribed (i.e. $\tilde{g}$ -type boundary conditions), are assigned equation number zero and are not included in the global ordering.

The stiffness matrix, $\mathbf{K}$, is made up of the two terms from the left-hand side of the integral equation. It is symmetric and positive definite, meaning that one can assemble only the upper triangular part, saving on both storage and operations. The matrix also possesses a band-profile structure which facilitates highly-efficient solution through the so-called ‘active-column’ skyline solver, in which zeros outside the profile are neither stored nor processed.

The right hand side, $\mathbf{f}$, is made up of three known parts, the body force term (f_i), the applied tractions (h_i) and the applied velocities (g_i).

In a finite element code such as CONMAN, it is most convenient to form the arrays $\mathbf{K}$ and $\mathbf{f}$ in an element-by-element fashion. In this regards, we may write:

$$\mathbf{K} = \mathbf{A}_{e=1}^{n_{el}}(\mathbf{k}^e), \quad \mathbf{f} = \mathbf{A}_{e=1}^{n_{el}}(\mathbf{f}^e), \quad (\text{B.46})$$

where $\mathbf{A}$ denotes an ‘assembly operator’ whose function is to add elemental contributions (namely $\mathbf{k}^e$ and $\mathbf{f}^e$) to the appropriate locations of $\mathbf{K}$ and $\mathbf{f}$, and n_{el} is the number of elements. It can be shown that the element arrays may be defined as follows:

$$\mathbf{k}^e = [k_{pq}^e], \quad \mathbf{f}^e = \{f_p^e\}, \quad 1 \leq p, q \leq n_{ee}, \quad (\text{B.47})$$

$$k_{pq}^e = \mathbf{e}_i^T k_{ab}^e \mathbf{e}_j, \quad p = n(a-1) + i, \quad q = n(b-1) + j, \quad 1 \leq a, b \leq n_{en}, \quad (\text{B.48})$$

$$\mathbf{k}_{ab}^e = \int_{\Omega^e} (\mathbf{B}_a^e)^T \mathbf{D}_\lambda \mathbf{B}_b^e d\Omega + \int_{\Omega^e} (\mathbf{B}_a^e)^T \mathbf{D}_\mu \mathbf{B}_b^e d\Omega \quad (\text{B.49})$$

$$f_p^e = \int_{\Omega^e} N_a^e f_i d\Omega + \int_{\Gamma_{he}} N_a^e h_i d\Gamma - \sum_{q=1}^{n_{ee}} k_{pq}^e g_q^e, \quad \Gamma_{he} = \Gamma_h \cap \Gamma^e \quad (\text{B.50})$$

$$g_q^e = g_j(\mathbf{x}_b^e) \quad \text{if } \mathbf{x}_b^e \in \Gamma_g, \quad (\text{B.51})$$

$$g_q^e = 0 \quad \text{if } \mathbf{x}_b^e \notin \Gamma_g. \quad (\text{B.52})$$

In the above, a superscript T denotes transpose; n_{en} is the number of element nodes; n_{ee} is the number of element equations; a and b are element (‘local’) node numbers; p and q are element (‘local’) equation numbers; N_a^e is the shape function associated with node a of the e th element; and $\mathbf{e}_i$ is the i th canonical basis vector of $\mathbb{R}^n$. Since CONMAN is a two-dimensional code:

$$\mathbf{e}_1 = \begin{Bmatrix} 1 \\ 0 \end{Bmatrix}, \quad \mathbf{e}_2 = \begin{Bmatrix} 0 \\ 1 \end{Bmatrix} \quad (\text{B.53})$$

The arrays $\mathbf{B}_a^e$, $\mathbf{D}_\lambda$ and $\mathbf{D}_\mu$ are given as:

$$\mathbf{B}_a^e = \begin{bmatrix} N_{a,1}^e & 0 \\ 0 & N_{a,2}^e \\ N_{a,2}^e & N_{a,1}^e \end{bmatrix} \quad \mathbf{D}_\lambda = \lambda \begin{bmatrix} 1 & 1 & 0 \\ 1 & 1 & 0 \\ 0 & 0 & 0 \end{bmatrix} \quad \mathbf{D}_\mu = \mu \begin{bmatrix} 2 & 0 & 0 \\ 0 & 2 & 0 \\ 0 & 0 & 1 \end{bmatrix} \quad (\text{B.54})$$

The reader is reminded that commas denote partial differentiation (e.g. $N_{a,1}^e = \partial N_a^e / \partial x_1$).

More details of the method and a formal error analysis are available in Hughes *et al.* (1979).

B.4 The Energy Equation

The energy equation is a simple advection-diffusion equation. The formal statement of the problem is: Find $T : \Omega \rightarrow \mathbb{R}$ such that:

$$\dot{T} + u_i T_{,i} = \kappa T_{,ii} + H \quad \text{on } \Omega \quad (\text{B.55})$$

$$T = b \quad \text{on } \Gamma_g \quad (\text{B.56})$$

$$T_{,j} n_j = q \quad \text{on } \Gamma_h \quad (\text{B.57})$$

Here, T is the temperature, $\dot{T}$ is the time derivative of temperature, u_i is the velocity in the i^{th} Cartesian direction, κ is the thermal diffusivity and H is an internal heat source.

The energy equation is solved using the Streamline Upwind Petrov-Galerkin (SUPG) method (Hughes & Brooks 1979). The Petrov-Galerkin (PG) schemes can be thought of as standard Galerkin methods in which the numerical under-diffusion inherent to such schemes is counterbalanced by adding an artificial diffusivity. The SUPG scheme further improves this counterbalance by removing crosswind diffusion - the PG weightings act only in the flow direction (i.e. it follows the streamline) hence the name (see Hughes & Brooks 1979, for further information).

As with the momentum equation, the preliminary step in the finite element formulation of the energy equation is to derive the weak form, stated as follows:

$$\int_{\Omega} (w + p) \dot{T} d\Omega = - \int_{\Omega} (w + p) (u_i T_{,i}) d\Omega - \kappa \int_{\Omega} w_{,i} T_{,i} d\Omega + \kappa \int_{\Gamma_h} w T_{,j} n_j d\Gamma_q + \int_{\Omega} (w + p) H \quad (\text{B.58})$$

In the above equation $T_{,i}$ is the gradient of temperature, w is the standard, continuous Galerkin weighting function and $(w + p)$ is the PG weighting function, with p the discontinuous, streamline upwind contribution. The modified weighting functions are applied on the time, internal heat source and advective terms only, where, in addition to correcting for under-diffusion, they remove the oscillations which result from the standard Galerkin formulation for an advection-dominated problem.

B.4.1 Predictor-Multicorrector Algorithm

The finite element discretization of Equation (B.58) leads to the following semi-discrete system of equations:

$$\mathbf{M}\dot{\mathbf{T}} + \mathbf{C}\mathbf{T} = \mathbf{F} \quad (\text{B.59})$$

where $\mathbf{M}$, $\mathbf{C}$, and $\mathbf{F}$ represent the ‘mass’ matrix, diffusion and convection matrix, and generalized force vector, respectively. $\mathbf{T}$ and $\dot{\mathbf{T}}$ represent vectors of nodal values of T and $T_{,t}$ respectively. Note that the SUPG formulation leads to a nonsymmetric mass matrix.

This initial value problem consists of finding a function $\mathbf{T} = \mathbf{T}(t)$ satisfying Equation (B.59) and the initial condition:

$$\mathbf{T}(0) = \mathbf{T}_0 \quad (\text{B.60})$$

where $\mathbf{T}_0$ is given.

The algorithm used to solve Equation (B.59) is a version of the predictor-multicorrector algorithms proposed by Hughes (1980). The solution procedure is as follows:

(1) Predictors:

$$\mathbf{T}_{n+1}^{(0)} = \mathbf{T}_n + \Delta t(1 - \alpha)\dot{\mathbf{T}}_n \quad (\text{B.61})$$

$$\dot{\mathbf{T}}_{n+1}^{(0)} = 0 \quad (\text{B.62})$$

(2) Solution:

$$\mathbf{R}_{n+1}^{(i)} = \mathbf{A}(\mathbf{f}_{n+1}^e - \mathbf{m}^e \dot{\mathbf{T}}_{n+1}^{e(i)} - \mathbf{c}^e \mathbf{T}_{n+1}^{e(i)}) \quad (\text{Residual Force}) \quad (\text{B.63})$$

$$\mathbf{M}\Delta\dot{\mathbf{T}}_{n+1}^{(i)} = \mathbf{R}_{n+1}^{(i)} \quad (\text{B.64})$$

(3) Correctors:

$$\mathbf{T}_{n+1}^{(i+1)} = \mathbf{T}_{n+1}^{(i)} + \Delta t\alpha\dot{\mathbf{T}}_{n+1}^{(i)} \quad (\text{B.65})$$

$$\dot{\mathbf{T}}_{n+1}^{(i+1)} = \dot{\mathbf{T}}_{n+1}^{(i)} + \Delta\dot{\mathbf{T}}_{n+1}^{(i)} \quad (\text{B.66})$$

(4) If an additional iteration is required, increment i , and go to step 2. If no additional iteration is needed, increment n , and go to step 1.

The notation is as follows: Δt is the time step; $\mathbf{T}_n$ and $\dot{\mathbf{T}}_n$ are the approximations to $\mathbf{T}(t_n)$ and $\dot{\mathbf{T}}(t_n)$ respectively; the superscript i is the iteration number (for the corrector); the superscript e denotes an element level matrix or vector and α is a convergence parameter. The symbol $\mathbf{A}$ represents an assembly operator which adds the element level force vectors to the global force vector. The element vectors $\mathbf{T}^e$ and $\dot{\mathbf{T}}^e$ include the prescribed boundary conditions in addition to the local values of the global $\mathbf{T}$ and $\dot{\mathbf{T}}$ vectors.

This algorithm is particularly powerful due to the residual (or ‘out-of-balance-force’) formulation and the ability to make multiple iterations within a time step. Additionally, the option of implicit or explicit treatment of specific terms is naturally contained within the algorithmic structure.

CONMAN uses an explicit time-stepping method. Consequently:

$$\mathbf{M} = \mathbf{M}^L \tag{B.67}$$

where $\mathbf{M}^L$ denotes the lumped mass matrix. The time step is dynamically selected through the Courant-Friedrichs-Levy criterion, corresponding to the largest step that can be taken explicitly and maintain stability. With the appropriate choice of variables, $\alpha = 0.5$ and two iterations, the method is second-order accurate.

Appendix C

Two Dimensional Mesh Generation and Adaptive Remeshing Procedure

C.1 Mesh Generation

The algorithmic procedure to be described for the mesh generation process is based upon the method originally proposed by Peraire et al. (1987). The advocated approach is regarded as a variant of the so-called ‘advancing front’ technique (George 1971, Lo 1985) with the distinctive feature that elements and nodes are generated simultaneously. This technique is capable of generating meshes that conform to an externally prescribed spatial distribution of element size. The ability to generate meshes that are locally stretched along prescribed directions is also included, leading to highly efficient definition of one-dimensional flow features. For simplicity, triangular elements are generated initially. These are subsequently combined or subdivided to form quadrilaterals, the elements utilized by CONMAN.




Figure C.1: An illustration of the advancing front technique. The figure shows different stages during the triangulation process.

C.1.1 Generation of the Initial Mesh

The underlying process in the advancing front technique is illustrated in Figure C.1 (see Peraire et al. 1987, 1990, for further discussion). The boundary of the domain is discretized first. Nodal points are placed on the boundary curves in such a way that the distance between them is as close as possible to the desired mesh spacing. Contiguous nodes on the boundary curves are joined by straight-line segments and assembled to form the initial generation front. At this stage, the triangulation loop begins. A side from the front is chosen and a triangle is generated that will have this selected side as one edge. In generating this new triangle an interior node may be created or an existing node in the front may be chosen. After generating the new element the front is conveniently updated in such a way that it always contains the sides that are available to form a new triangle. The generation is complete when no sides are left in the front.




Figure C.2: The definition of the mesh parameters α , s and δ .

C.1.2 Mesh Control: The Background Mesh

The inclusion of adequate mesh control is a key ingredient in ensuring the generation of a mesh of the desired form. Control over the characteristics is obtained by the specification of a spatial distribution of mesh parameters by means of a background mesh. The background mesh is used for interpolation purposes only and is made up of triangles.

To control the elements generated, the user defines the node spacing, δ , the value of a stretching parameter, s , and a direction of stretching α . The generated elements will then have a typical length $s\delta$ in the direction parallel to α , and a typical length δ in the direction normal to α (Figure C.2). Thus, at each node on the background grid, the nodal values of δ , s , and α must be specified. These values are dependent upon the solution gradients and curvatures yielded by the error indication process (described further in Sections C.2.1 and C.2.2). Local values of these quantities are then obtained during the generation process by cubic interpolation, over the triangles of the background grid, between the specified nodal values (see Nielson 1979, El Hachemi et al. 2003, for a detailed description of this interpolation process).




Figure C.3: Orientation of the boundary. The domain of interest is shaded.

In general, for the initial mesh, the location of one-dimensional features is not known. Consequently a value of $s = 1$ (i.e. no stretching) is specified. The node spacing δ is also defined to be uniform, although a variation of δ can be achieved (by suitable construction of the background grid) if it is apparent that increased mesh resolution is required in certain regions of the flow domain. Note that if δ is required to be uniform initially and no stretching is to be specified, then the background grid need only consist of a single element, which completely covers the solution domain.

C.1.3 Boundary Representation

The boundary of the solution domain is represented by closed loops of oriented piecewise cubic spline curves. For simply connected domains, these boundary curves are orientated in a counter-clockwise sense, while for multi-connected regions the exterior boundary curves are given a counter-clockwise orientation and all interior boundary curves are orientated in a clockwise sense (Figure C.3). This means that, as the boundary curve is traversed, the region to be meshed always lies to the left.

When these boundary curves are discretized, the boundary edges forming the initial front are




Figure C.4: Generation of a new triangle.

orientated in the same fashion. In this study, the orientation of a boundary edge is defined by the order in which the two nodes of the edge are listed in the front. The orientation of the edge is important as it identifies the area of the plane in which a valid triangle can be created using that edge as a base.

C.1.4 Triangle Element Generation

The generation of a regular triangular element of size δ is illustrated in Figure C.4. The process involves the following steps:

1. Select an edge AB from the generation front.
2. Using the orientation of the edge, determine the position of a point C_1 , which lies at a distance δ from A and B.
3. Determine all points in the front that lie inside a circle of radius δ and center at C_1 . Let these points be denoted by Q_i , where the subscript i varies between 0 and the

number of ‘front’ points inside the circle ($i = 2$ in Figure C.4).

4. Determine the positions of the equally spaced points C_2 , C_3 , C_4 and C_5 on the line joining C_1 and the mid-side (M) point of AB .
5. Form a list containing all the points determined in step 3 as well as points C_1 , C_2 , C_3 , C_4 and C_5 . The points in this list will then be ordered according to their distance from the point C_1 .
6. Create an element with nodes A , B and the first point in the list which satisfies the mesh consistency requirement, i.e. the two newly created edges do not cross any of the existing edges in the front.
7. Update the front by removing the edge AB , and adding the appropriate number of new edges with the correct orientation.

C.1.5 Quadrilateral Element Generation using an Existing Triangular Mesh

Since CONMAN can only handle quadrilateral grids, the triangular meshes generated by the procedures described above must be altered into quadrilaterals. This can be done by:

1. Combining two triangles together to form a quadrilateral. However, this method comes unstuck when an odd number of triangular elements exist in the original mesh.
2. Generating mid-side nodes on triangles and one in the center, and interconnecting them, generating three quadrilaterals.

In this study, a combination of the two methods is employed (Figure C.5):

1. The triangles of the mesh are combined in pairs to give quadrilaterals.
2. Since there may be several triangles remaining, all elements, triangular and quadrilateral, are split into quadrilaterals by placing new nodes at the mid-sides, and one in the middle of each element. This ensures the generation of an all-quadrilateral mesh.

It should be pointed out that in order to produce a mesh with desired element size, δ , the spacing of the triangles generated originally must be 2δ .




Figure C.5: Agglomeration of triangles.

C.1.6 Mesh Quality Enhancement

Three post-processing procedures are applied to enhance mesh quality. These procedures do not alter the total number of points or elements in the mesh:

1. *Diagonal Swapping*: This technique changes the connectivities among nodes in the mesh without altering their position. This process requires a loop over all the element sides excluding those sides on the boundary. For each side common to the triangles ACD and BCD (see Figure C.6), one considers the possibility of swapping CD by AB, thus replacing the two triangles ACD and BCD by the triangles ABC and ABD. The swapping is performed if a prescribed regularity criterion is better satisfied by the new configuration than by the existing one. In our implementation, the swapping operation is performed if the minimum angle (formed by element edges at each element node) occurring in the new configuration is larger than in the original one.
2. *Mesh Smoothing*: This alters the position of the interior nodes without changing the topology of the mesh. The element sides are considered as springs of stiffness proportional to the length of the side. The nodes are moved until the spring system is in equilibrium. The equilibrium positions are found by iteration. Each iteration amounts to performing a loop over the interior points and moving their coordinates to coincide with those of the centroid of the neighboring points. Usually three to five




Figure C.6: Diagonal swapping procedure.

iterations are performed (Hassan & Probert 1999).

3. *Bandwidth Reduction*: A variant of the 'Cuthill-McKee' algorithm (an algorithm to reduce the bandwidth of sparse symmetric matrices) is utilized to renumber elements and nodes (Cuthill & McKee 1969). This process speeds up the calculation and reduces memory requirements, since the resulting solution matrices are far more compact, with reduced bandwidth.

C.2 The Error Indicator and Adaptive Strategy

Having obtained an initial approximation to a solution for a given problem, one can improve the accuracy of this solution by adaptively refining the mesh. In this study, mesh adaptation is achieved by using the computed solution to determine 'optimum' nodal values for δ , s , and α . The mesh is then regenerated with the initial computational mesh acting as a background grid.

To determine the values for the mesh parameters, it is necessary to use the initial solution to give some indication of the error magnitude and direction. A certain 'key' variable must be

identified and then the error indication process can be performed in terms of this variable. In this study, the error indicator is based on the temperature variable, T , in purely thermal convective simulations, and a combination of T and C , the composition, in thermo-chemical simulations. Of course, other variables (e.g. pressure) or any combination of variables (e.g. temperature and velocity, Nithiarasu 2000) can be chosen, depending upon the nature of the problem under investigation.

The construction of the error indicator can take various forms depending upon the nature of the problem. It is obvious from the large number of publications available on error estimation and adaptivity (Lohner et al. 1986, Zienkiewicz & Zhu 1987, Lohner 1995, Hassan et al. 1995, Fortin et al. 1996) that research in this area of computational mechanics remains very active. However, as has been pointed out by Nithiarasu & Zienkiewicz (2000), most of the well-known literature on error estimates deals with self-adjoint problems (e.g. Zienkiewicz & Zhu 1987). Fluid mechanics problems, which involve non self-adjoint operators, are more difficult concepts, and traditional methods such as the 'energy norm' are not always suitable for measuring the error. For this reason, most of the work in fluids utilizes local indicators, such as the local interpolation error, to refine the grid without specifying a total error.

Error indicators based upon the interpolation theory make the assumptions that:

1. The nodal error is zero.
2. The solution is smooth.

This allows one to *approximate* the elemental error by a derivative one order higher than the element shape function. We make use of this approach to refine the grid, by considering the second derivatives (curvatures) of T and C . Note that for the remainder of this appendix we will restrict our discussion to the solution variable ϕ , rather than refer to T and C explicitly.

C.2.1 The Error Indicator

Consider a one-dimensional situation in which the exact values of the key variable ϕ are approximated by a piecewise linear function $\hat{\phi}$. The error E is then defined as:

$$E = \phi(x) - \hat{\phi}(x) \quad (C.1)$$

If the exact solution is a linear function of x , this error will vanish, as the approximation has been obtained using piecewise linear finite element shape functions. To a first order of approximation, the error E can be evaluated as the difference between a quadratic finite element solution $\hat{\phi}$ and the linear computed solution. To obtain a piecewise quadratic approximation, one could obviously solve a new problem using quadratic shape functions. This, however, would be costly and an alternative approach for estimating a quadratic approximation from the linear finite element solution can be employed. Assuming that the nodal values of the quadratic and the linear approximations coincide i.e. that the nodal values of E are zero, a quadratic solution can be constructed on each element, once the value of the second derivative is known (assuming that second derivatives are constant over each element).

The variation of the error E within an element e is then expressed as:

$$E_e = \frac{1}{2}\zeta(h_e - \zeta)\frac{\partial^2 \hat{\phi}}{\partial x^2} \quad (C.2)$$

where ζ denotes a local element coordinate and h_e denotes the element length (Peraire *et al.* 1987). The root mean square value E_e^{RMS} of this error over the element is computed as:

$$E_e^{RMS} = \left\{ \int_0^{h_e} \frac{E_e^2}{h_e} d\zeta \right\}^{1/2} = \frac{1}{\sqrt{120}} h_e^2 \left| \frac{\partial^2 \hat{\phi}}{\partial x^2} \right|_e \quad (C.3)$$

where $||$ denotes absolute value. Several previous studies (Demkowicz *et al.* 1984, Peraire *et al.* 1987, Nithiarasu 2000) have demonstrated that equi-distribution of the element error leads to an optimal mesh and in what follows we employ the same criterion. This

requirement implies that:

$$h^2 \left| \frac{\partial^2 \hat{\phi}}{\partial x^2} \right| = C \quad (\text{C.4})$$

where C denotes a positive constant. Finally, the requirement of Equation (C.4) suggests that the optimal spacing δ on the new adapted mesh should be computed according to:

$$\delta^2 \left| \frac{\partial^2 \hat{\phi}}{\partial x^2} \right| = C \quad (\text{C.5})$$

Equation (C.5) can be directly extended to the two-dimensional case by writing the quadratic form:

$$\delta_\beta^2 (m_{ij} \beta_i \beta_j) = C \quad (\text{C.6})$$

where β is an arbitrary unit vector, δ_β is the spacing along the direction of β , and m_{ij} are the components of a 2×2 symmetric matrix, $\mathbf{m}$, of second derivatives defined by:

$$m_{ij} = \frac{\partial^2 \hat{\phi}}{\partial x_i \partial x_j} \quad (\text{C.7})$$

These derivatives are computed at each node of the current mesh by using the two-dimensional equivalent of the variational recovery procedure. This procedure allows one to recover the nodal values of second derivatives from the elemental values of the first derivatives of $\hat{\phi}$ - refer to Appendix D for a detailed analysis of this procedure.

C.2.2 Adaptive Remeshing

The basic concept behind the adaptive remeshing technique is to use the computed solution to provide information on the spatial distribution of the mesh parameters. This information will be used by the mesh generator to generate a new adapted mesh in those areas where the values of the optimal mesh parameters differ from the values of the current mesh parameters by greater than a user prescribed amount, err_max (set as 0.5% in this study).

The optimal values for the mesh parameters are calculated at each node of the current mesh. The directions $\alpha_i; i = 1, 2$ are taken to be the principal directions of the matrix $\mathbf{m}$. The corresponding mesh spacings are computed from the eigenvalues λ_i of $\mathbf{m}$, as:

$$\delta_i = \sqrt{\frac{C}{\lambda_i}}; i = 1, 2 \quad (\text{C.8})$$

The spatial distribution of the mesh parameters is defined when a value is specified for the constant C . The total number of elements in the adapted mesh will depend upon the choice of this constant. The magnitude of the stretching parameter, s , at node n , is simply defined as the ratio between the two spacings:

$$s_n = \sqrt{\frac{|\delta_{1n}|}{|\delta_{2n}|}} \quad (\text{C.9})$$

where δ_{1n} is the spacing in principal direction 1, and δ_{2n} is the spacing in principal direction 2.

In the practical implementation of this method, two threshold values are used: a minimum spacing δ_{min} , and a maximum spacing δ_{max} , with:

$$\delta_{min} \leq \delta_i \leq \delta_{max}; i = 1, 2 \quad (\text{C.10})$$

It is apparent that in regions of uniform flow, the computed values of δ_n will be very large. Consequently, the user must specify a maximum allowable value, δ_{max} , for the local spacing on the new mesh. Then, if δ_n is such that $\delta_n \geq \delta_{max}$, the value of δ_n is set to δ_{max} . Similarly, the user prescribes a maximum allowable stretching ratio on the new mesh.

The new mesh is generated according to the computed distribution of mesh parameters. The original solution is then transferred onto the new mesh using cubic interpolation (Nielson 1979, El Hachemi et al. 2003) and the solution procedure continues on the new mesh. It should be noted that the increase in definition of flow features is achieved by decreasing the

value of δ_{min} . The value of δ_{min} is therefore the major parameter governing the number of elements in the new mesh.

Appendix D

Calculation of Nodal Gradients and Curvatures

The discrete, finite element solution provides the values of ϕ^h in terms of the nodal values $\bar{\phi}$ as:

$$\phi^h = \mathbf{N}\bar{\phi} \quad (\text{D.1})$$

where $\mathbf{N}$ are the appropriate shape functions used. The nodal values of the second derivatives, $\partial^2\phi^h/\partial x^2$, can be obtained by using a similar approximation:

$$\frac{\partial^2\phi^h}{\partial x^2} = \mathbf{N}\left\{\frac{\partial^2\bar{\phi}}{\partial x^2}\right\} \quad (\text{D.2})$$

with similar expressions for $\frac{\partial^2\phi^h}{\partial y^2}$. The projection:

$$\int_{\Omega} \mathbf{N}^T \left[\mathbf{N} \left\{ \frac{\partial^2\bar{\phi}}{\partial x^2} \right\} - \frac{\partial^2\mathbf{N}}{\partial x^2} \bar{\phi} \right] d\Omega = 0 \quad (\text{D.3})$$

can be used to determine the nodal curvatures. Thus:

$$\frac{\partial^2\bar{\phi}}{\partial x^2} = \mathbf{M}^{-1} \int_{\Omega} \left\{ \frac{\partial\mathbf{N}^T}{\partial x} \frac{\partial\mathbf{N}}{\partial x} \right\} d\Omega \bar{\phi} \quad (\text{D.4})$$

where:

$$\mathbf{M} = \int_{\Omega} \mathbf{N}^T \mathbf{N} d\Omega \quad (\text{D.5})$$

is the well-known mass matrix, which is lumped for convenience. The contribution at any node thus involves only the elements surrounding it. Similar expressions can be written for $\partial^2 \bar{\phi} / \partial y^2$ and $\partial^2 \bar{\phi} / \partial x \partial y$.

Appendix E

TERRA

E.1 Overview & History

TERRA is a well established 3-D finite element mantle dynamics code that was first developed by Baumgardner (1983). The code solves for momentum and energy balance at infinite Prandtl number in a spherical shell, with the inner radius being that of Earth's outer core and the outer radius corresponding to Earth's surface (although these radii are not fixed). The version of TERRA described in this chapter corresponds to a modified version of the original code. Several revisions have played their part in rendering TERRA arguably the most functional and well known mantle dynamics code available today. The most notable modifications and additions to the code are listed below:

1. The code was parallelized by Bunge & Baumgardner (1995) in the most significant alteration since its creation. By utilizing a number of processors, the inherent limitations of single processor machines are overcome. Mantle simulations can therefore be examined with unprecedented accuracy at far more realistic resolutions.
2. Yang (1997) modified the code considerably to allow for variable viscosity simulations (both temperature and pressure dependent). With viscosity playing such an important role in mantle convection, this was an essential upgrade.

3. Plate like surface velocities were incorporated by Bunge et al. (1998). Although not representing ‘true’ plate tectonic behavior (lithospheric deformation is not coupled to lithosphere mechanics), these imposed velocities have proved highly useful in many mantle circulation studies.
4. Tracer particle subroutines were added to the code, allowing investigations into thermochemical phenomena within Earth’s mantle (Stegman et al. 2003). Such a capability will no doubt prove useful in the future in resolving various geochemical and geophysical contradictions.
5. Oldham (2004) created a ‘marker-chain’ version of the code allowing for layered mantle convection simulations. Although recent evidence points towards whole mantle convection as the most likely geometry of mantle convection, various layered models are still being supported and this code will doubtless prove useful in testing these in the future.

E.2 The Governing Equations

As noted in Chapter 1, Earth’s mantle is solid. However, over large timescales it deforms slowly through processes such as dislocation and diffusion creep. As a consequence, motion within Earth’s mantle can be described by the equations governing fluid dynamics. However, since the mantle has an extremely large viscosity ($\approx 10^{21}$ Pa s), the equations governing mantle convection are somewhat different to those governing the more ‘*typical*’ fluid mechanics problems:

- The large viscosity of Earth’s mantle makes the Prandtl Number (Pr), the ratio between viscous and inertial forces, of the order 10^{24} . Accordingly, inertial terms in the momentum equation can be ignored.
- The Ekman number (i.e. the ratio between viscous and Coriolis forces) is of the order 10^9 , since the velocity of convection within the mantle is so small. As a consequence, the Coriolis force can be neglected.

- The centrifugal force, which causes the ratio of major radius to minor radius of Earth to depart from unity by only one part in three hundred, is proportional to the square of the velocity. Consequently, it is even smaller than the Coriolis force and it is also ignored.

This mantle convection problem is formulated in terms of the conservation equations of linear momentum, mass and energy, inside a spherical shell with appropriate boundary conditions. Unless stated otherwise, for the results presented in this thesis, a linear and isotropic constitutive law is assumed so that the mantle is treated as an infinite Prandtl number Newtonian fluid (although TERRA is more flexible, allowing investigations into various rheologies). Under these assumptions, the following vector equations describe the local behavior:

$$\nabla \cdot \tau - \nabla p + \rho g = 0 \quad (\text{E.1})$$

$$\frac{\partial \rho}{\partial t} = -\nabla \cdot (\rho \mathbf{u}) \quad (\text{E.2})$$

$$\frac{\partial T}{\partial t} = -\nabla \cdot (T \mathbf{u}) - (\gamma - 1) T \nabla \cdot \mathbf{u} + \frac{[\tau : \nabla \mathbf{u} + \nabla \cdot (k \nabla T) + H]}{\rho c_V} \quad (\text{E.3})$$

where

$$\tau = \mu \left[\nabla \mathbf{u} + (\nabla \mathbf{u})^T - \frac{2I(\nabla \cdot \mathbf{u})}{3} \right] \quad (\text{E.4})$$

and

$$p = p(\rho, T) \quad (\text{E.5})$$

In the above equations p denotes pressure, ρ density, g gravitational acceleration, τ deviatoric stress, $\mathbf{u}$ fluid velocity, T absolute temperature, γ the Grüneisen parameter, k thermal conductivity, H volumetric radiogenic heat production, c_V specific heat at constant volume and μ dynamic viscosity, which can be a function of temperature and pressure. Equation (E.1) is the force-balance equation, describing the balance between pressure gradient, buoyancy forces and viscous forces. Equation (E.2) expresses the conservation of mass, while

Equation (E.3) represents the conservation of energy in terms of the absolute temperature. Equation (E.4) is the linear constitutive law, relating the stress field to the velocity, with Equation (E.5) corresponding to the equation of state as a suitable function of density and temperature.

For completeness, Equations (E.1)–(E.4) are now expressed in the tensor notation. This allows for a more transparent derivation of the finite element matrix equations in the following sections.

$$\frac{\partial \tau_{lm}}{\partial x_m} - \frac{\partial p}{\partial x_l} + \rho g_l = 0 \quad (\text{E.6})$$

$$\frac{\partial \rho}{\partial t} = -\frac{\partial(\rho u_l)}{\partial x_l} \quad (\text{E.7})$$

$$\frac{\partial T}{\partial t} = -\frac{\partial(T u_l)}{\partial x_l} - (\gamma - 1) T \frac{\partial u_l}{\partial x_l} + \frac{1}{\rho c_V} \left[\tau_{lm} \frac{\partial u_m}{\partial x_l} + \frac{\partial}{\partial x_l} \left(k \frac{\partial T}{\partial x_l} \right) + H \right] \quad (\text{E.8})$$

$$\tau_{lm} = \mu \left(\frac{\partial u_l}{\partial x_m} + \frac{\partial u_m}{\partial x_l} - \frac{2}{3} \delta_{lm} \frac{\partial u_k}{\partial x_k} \right) \quad (\text{E.9})$$

For incompressible cases, the mass conservation equation (E.2 / E.7) is replaced by the incompressibility approximation:

$$\nabla \cdot \mathbf{u} = 0 \quad \text{or} \quad \frac{\partial u_m}{\partial x_m} = 0$$

The Boussinesq formulation, where density variations are neglected everywhere except in the buoyancy term, is also employed in such cases.

The vigor of convection is described by the Rayleigh number, defined in bottom heated cases as:

$$Ra_T = \frac{\rho g \beta \Delta T h^3}{\kappa \mu} \quad (\text{E.10})$$

where ΔT is the temperature difference between the upper (Earth's surface) and lower (core-mantle boundary) boundaries of the spherical shell, h is the distance between these

boundaries, β is the coefficient of thermal expansion and κ is the thermal diffusivity, given by $k/(\rho c_P)$. The specific heat at constant pressure, c_P , is related to c_V by:

$$c_P = (1 + \gamma\beta T) c_V \quad (\text{E.11})$$

For internally heated cases, the basally heated Rayleigh number, Ra_T , is replaced by the internally heated Rayleigh number, Ra_H , given by:

$$Ra_H = \frac{\rho^2 g \beta H h^5}{k \kappa \mu} \quad (\text{E.12})$$

The equation set is now completely defined. The strategy for solving this set of equations is:

1. Compute the pressure field from the density and temperature fields via the equation of state (E.5).
2. Solve Equations (E.1) and (E.2) simultaneously for $\mathbf{u}$ and p .
3. Solve Equation (E.3) for the rate of change of T .
4. Take a time step and update the density and temperature fields.

The remainder of this appendix will focus on the numerical methods employed in solving these equations. The underlying numerical grid is described first (Section E.3). This is followed by a derivation of the discretized Galerkin finite element equations (Section E.4). The transformation of these equations into computational operators, in terms of the finite element shape functions, is then briefly described in Section E.5, together with the relevant solution strategies (Section E.6). The energy equation is solved by means of a finite difference method, in conjunction with the finite volume method, using the aforementioned finite element computational grid. A description of this procedure is included in Section E.7 along with an explanation of the time stepping strategies employed. The appendix ends

with an outline of the parallelization within TERRA (Section E.8) and a brief overview of the visualization package utilized during post-processing (Section E.9).

It is important to point out that the description of TERRA given here does not include the modifications made during my PhD work. These are described in Chapter 4 of this thesis. This appendix merely aims to provide a basic introduction to the code, focussing on the aspects of TERRA that are central to my PhD studies. In areas where material is covered only briefly, the reader will be referred to more detailed accounts in selected references.

E.3 The Computational Grid

The computational grid employed by TERRA is described in detail by Baumgardner (1983) and Baumgardner & Frederickson (1985). A brief summary is presented here, although a more detailed account can be found in these references.

An integral part of the finite element method involves splitting the computational domain into elements. As had been noted previously, this subdivision is termed *meshing* or *discretization*. In TERRA, the thick spherical shell is discretized by an icosahedral grid. By projecting the regular icosahedron onto a sphere, the spherical surface can be divided into twenty identical spherical triangles (or ten identical diamonds, each of which contains one of the ten triangles surrounding each of the two poles). Each triangle can subsequently be subdivided into four triangles by construction of great circle arcs between triangle side midpoints. This refinement process can be repeated, as shown in Figure E.1, to yield an almost uniform triangulation of the sphere at any desired resolution. Unlike conventional discretization methods on the sphere, since the icosahedral scheme achieves an almost uniform triangulation of the spherical surface, it avoids the ‘pole problem’ of traditional latitude-longitude grids. A larger time step can therefore be exploited during numerical simulations,




Figure E.1: (a) Mesh produced by projection of the regular icosahedron onto the sphere. (b) - (d) represent successive mesh refinements obtained by connecting midpoints of triangle sides with great circle arcs. By studying the grid it can be seen that most nodes (i.e. points at the corners of spherical triangles) are surrounded by six triangles. However, twelve nodes, at the vertices of the original ten diamonds, are enclosed by only five triangles. Reproduced with permission from David Oldham (Oldham 2004).

since elements are all of a similar dimension (as a consequence of the Courant-Friedrichs-Levy stability condition).

Within the TERRA code, refinements to the grid are referenced by ‘the number of grid intervals along an icosahedral diamond edge’. This value is labeled mt . The value of mt defines the number of data points present in the TERRA grid and, hence, its resolution. The number of data points on a spherical surface is given by $10(mt^2 + 1) + 2$, since there are 10 icosahedral diamonds on each surface and two polar nodes.

The grid is extended radially by placing several of these spherical shells above one and other, generating a mesh of triangular prisms with spherical ends (see Figure E.2). Yang (1997) states that for Earth’s mantle, the best performing grid is obtained when nr , the number of radial layers, is set to equal $mt/2$. Unless otherwise stated, such a configuration is adopted throughout this thesis. Having extended the grid in the radial direction, the total number of nodes in the spherical shell is given by: $(nr + 1)[10(mt^2 + 1) + 2]$. Note that there are $nr + 1$ radial layers.

E.3.1 Indexing Conventions Employed by Terra

The icosahedral grid leads to a convenient data structure for numerical simulations. This data structure is described here.

The original grid is divided into twenty equilateral triangles. These twenty triangles are grouped into pairs to form ten diamonds, each of which surrounds one of the two pentagonal vertices, the poles. These diamonds are indexed as shown in Figure E.3. Within the TERRA program, the diamond number is referred to as id .

The data points within each diamond (or more accurately each subdomain - see parallelization section) are referenced by a pair of indices. The values of these indices are derived




Figure E.2: Indexing convention for radial discretization of a spherical shell. Reproduced with permission from John Baumgardner (Baumgardner 1983).




Figure E.3: Indexing convention for the ten icosahedral diamonds: (a) Viewed from the Equator (b) Viewed from the North Pole (NP) (c) Viewed from the South Pole (SP) and (d) With diamonds projected onto a flat plane. Reproduced with permission from David Oldham (Oldham 2004).




Figure E.4: The indexing convention for diamonds in an $mt = 8$ case. (a) Northern hemisphere diamonds (b) Southern Hemisphere diamonds. Hexagons represent data points. Solid grey hexagons represent data points that are owned by the diamond while transparent grey hexagons are data points that, while they do form part of this diamond, are on the edge of the diamond and so are considered to be owned by the adjacent diamond. Reproduced with permission from David Oldham (Oldham 2004).

from axes which run along the diamond edges. These axes have their origin at the poles, with nodes labeled $i1$ and $i2$, as shown in Figure E.4.

Finally, the radial layers are indexed by ir , with values ranging from 1 (at the surface) to $nr + 1$ (at the core-mantle boundary). Any point within the computational domain can, therefore, be described by four indices: $i1$, $i2$, id and ir .

E.4 The Galerkin Finite Element Formulation

The discretization of the differential equations (E.1 - E.4) using the Galerkin finite element formulation is described next. Substituting Equation (E.4) into Equation (E.1), yields:

$$\mathbf{A}\mathbf{u} - \nabla p = \mathbf{b} \quad (\text{E.13})$$

where $\mathbf{A}$ is the differential operator associated with divergence of stress and $\mathbf{b}$ is the buoyancy force vector. The equations we wish to solve (for the approximate solutions $\tilde{\mathbf{u}}$ and $\tilde{p}$) are, therefore, the Momentum equation:

$$\mathbf{A}\tilde{\mathbf{u}} - \nabla\tilde{p} \approx \mathbf{b} \quad (\text{E.14})$$

and the mass conservation equation, given by either:

$$\nabla \cdot \tilde{\mathbf{u}} \approx 0 \quad (\text{E.15})$$

for incompressible fluids or:

$$\nabla \cdot (\rho\tilde{\mathbf{u}}) \approx 0 \quad (\text{E.16})$$

for compressible fluids. For simplicity, the remainder of this appendix will consider only incompressible fluids.

As explained in Appendix A, the Galerkin finite element method is one of several weighted residual formulations which work by minimizing the residuals of the above equations in the volume integral sense. The major advantage of the Galerkin formulation is that the finite element basis/shape functions themselves become the weighting functions. Equations (E.14) and (E.15) therefore become (in their Galerkin weighted residual form):

$$\int N_i (\mathbf{A}\tilde{\mathbf{u}} - \nabla\tilde{p} - \mathbf{b}) dV = 0 \quad (\text{E.17})$$

$$\int N_i (\nabla \cdot \tilde{\mathbf{u}}) dV = 0 \quad (\text{E.18})$$

where $\int(\cdot)dV$ represents an integration over the entire computational domain, i.e. the spherical shell. The next stage involves defining Equations (E.17) and (E.18) in terms of the finite element shape functions.

E.4.1 Matrix Equation Derivation

Transferring to tensor notation and remembering that in the finite element formulation, $\mathbf{u}$ and p can be related to their approximations $\tilde{\mathbf{u}}$ and $\tilde{p}$ by Equation (E.19) and Equation (E.20) respectively (note that l in u_l represents the l^{th} direction), the process of equation discretization is now described.

$$\mathbf{u}_l = N_i \tilde{\mathbf{u}}_{il} \quad (\text{E.19})$$

$$p = N_i \tilde{p}_i \quad (\text{E.20})$$

For simplicity, tildes will be dropped from $\tilde{\mathbf{u}}$ and $\tilde{p}$, with $\mathbf{u}$ and p denoting numerically determined approximations for the remainder of this section.

Focussing our attention on the divergence of stress term, $\mathbf{A}\mathbf{u}$, in its present form, the statement requires continuity of first derivatives of the shape functions if infinite values are to be avoided (see Appendix A.4.4). Integration by parts relaxes this requirement and leads to a weak form of the weighted residual statement. A natural boundary condition is also engendered by application of the Green-Gauss divergence theorem:

$$\int N_i \frac{\partial \tau_{lm}}{\partial x_m} dV = - \int \frac{\partial N_i}{\partial x_m} \tau_{lm} dV + \oint N_i \tau_{lm} \hat{n}_m dS \quad (\text{E.21})$$

In the above equation, $\oint(\cdot)dS$ represents an integration over the surface of the domain V , and $\hat{n}$ represents the unit outward normal vector, which becomes the radial unit vector, $\hat{\mathbf{r}}$ at the outer surface of the spherical shell, and $-\hat{\mathbf{r}}$ at the inner surface.

Exploiting Equations (E.9) and (E.19), the volume integral in Equation (E.21) becomes:

$$\begin{aligned} - \int \frac{\partial N_i}{\partial x_m} \tau_{lm} dV &= - \int \mu \frac{\partial N_i}{\partial x_m} \left(\frac{\partial u_l}{\partial x_m} + \frac{\partial u_m}{\partial x_l} - \frac{2}{3} \delta_{lm} \frac{\partial u_k}{\partial x_k} \right) dV \\ &= A'_{iljm} u_{jm} \end{aligned} \quad (\text{E.22})$$

where A' , the volume integral part of the total differential operator A , is defined as:

$$A'_{iljm} = - \int \mu \left(\delta_{lm} \frac{\partial N_i}{\partial x_k} \frac{\partial N_j}{\partial x_k} + \frac{\partial N_i}{\partial x_m} \frac{\partial N_j}{\partial x_l} - \frac{2}{3} \frac{\partial N_i}{\partial x_l} \frac{\partial N_j}{\partial x_m} \right) dV$$

$$(i, j = 1 \dots n; \quad l, m = 1, 2, 3) \quad (\text{E.23})$$

The third term inside the parentheses can be dropped in the case of incompressible fluids, and when incompressible fluids with constant viscosity are considered, one can include only the first term (the Laplacian operator), since $\mathbf{A}\mathbf{u}$ can now be written $\mu \nabla^2 \mathbf{u}$. For such incompressible, isoviscous cases, it is sufficient to consider the scalar operator:

$$S'_{ij} = - \int \mu \frac{\partial N_i}{\partial x_k} \frac{\partial N_j}{\partial x_k} dV \quad (\text{E.24})$$

as in the original version of the code (Baumgardner, 1983). However, when the fluid is compressible, or has variable viscosity, the full operator A' must be considered.

The integrand of the surface integral in Equation (E.21) can be significantly simplified with an application of the shear stress free boundary condition. This allows us to write:

$$\tau_{lm} \hat{r}_m = \tau_{rr} \hat{r}_l \quad (\text{E.25})$$

at a spherical surface. A full explanation of the mathematics behind this assumption can be found in Yang (1997). Combining the above equation and the identity:

$$\tau_{rr} = \mu \left(2 \frac{\partial u_r}{\partial r} - \frac{2}{3} \nabla \cdot \mathbf{u} \right) \quad (\text{E.26})$$

yields:

$$\oint N_i \tau_{lm} \hat{r}_m dS = A''_{iljm} u_{jm} \quad (\text{E.27})$$

where:

$$A''_{iljm} = \int_{s_1} \mu N_i \left(2 \frac{\partial N_j}{\partial r} \hat{r}_m - \frac{2}{3} \frac{\partial N_j}{\partial x_m} \right) \hat{r}_l dS_1$$

$$- \int_{s_2} \mu N_i \left(2 \frac{\partial N_j}{\partial r} \hat{r}_m - \frac{2}{3} \frac{\partial N_j}{\partial x_m} \right) \hat{r}_l dS_2$$

$$(i, j = 1 \dots n; \quad l, m = 1, 2, 3) \quad (\text{E.28})$$

Here, $\int_{s_1}(\cdot)dS_1$ and $\int_{s_2}(\cdot)dS_2$ represent integrations over the upper surface, S_1 , and lower surface, S_2 , respectively. As with the volume integral, the terms containing the factor $\frac{2}{3}$ can be ignored for incompressible fluids. Additionally, the surface integrals vanish when node i is not on either surface or when node j is not located on either a boundary or the layer adjacent to it. The total forward operator, A , is given by:

$$A_{iljm} = A'_{iljm} + A''_{iljm} \quad (\text{E.29})$$

The remaining integrals in Equations (E.17) and (E.18) can be transformed with relative ease. The pressure gradient term can be written as:

$$\int N_i \frac{\partial p}{\partial x_l} dV = \int N_i \frac{\partial N_j}{\partial x_l} dV p_j = G_{ilj} p_j \quad (\text{E.30})$$

with the gradient operator matrix G , defined as:

$$G_{ilj} = \int N_i \frac{\partial N_j}{\partial x_l} dV \quad (i, j = 1 \dots n; \quad l = 1, 2, 3). \quad (\text{E.31})$$

The integral containing the buoyancy force term can be written as:

$$\begin{aligned} f_{il} &= \int N_i b_l dV = - \int N_i \rho_0 \beta [T - T_0(r)] g \hat{r}_l dV \\ &\quad (i = 1 \dots n; \quad l = 1, 2, 3) \end{aligned} \quad (\text{E.32})$$

while the integral for the divergence of velocity in Equation (E.18) is given by:

$$\begin{aligned} \int N_i \nabla \cdot \mathbf{u} dV &= \int_{S_1} N_i u_r dS_1 - \int_{S_2} N_i u_r dS_2 - \int N_j \frac{\partial N_i}{\partial x_l} dV u_{jl} \\ &= -(\mathbf{u}^T G)_i \quad (i = 1 \dots n). \end{aligned} \quad (\text{E.33})$$

The surface integrals vanish since there is no radial velocity component on boundaries.

We finally arrive at the discrete form of the Navier-Stokes and incompressibility equations, expressed in discretized field variables:

$$\mathbf{A} \mathbf{u} - G p = \mathbf{f} \quad (\text{E.34})$$

$$G^T \mathbf{u} = 0 \quad (\text{E.35})$$

For ease of notation, Equation (E.34) can be simplified further to:

$$\mathbf{A} \mathbf{u} = \mathbf{f} \quad (\text{E.36})$$

where the right hand side term, $\mathbf{f}$, includes the term Gp . These equations are solved for a set of $4n$ unknowns, with n being the number of nodes and the scalar p and vector $\mathbf{u}$ representing the four independent variables.

E.5 Computational Operators

The previous section illustrates how to express the governing differential equations in terms of finite element shape functions and their derivatives. In this section, we briefly look into how these expressions are transformed into ones that can be readily calculated on the computational grid introduced in section E.3.

E.5.1 The Finite Element Shape Functions

For this section, nodes are referenced using two separate indices, the first denoting its radial position, with the second denoting its tangential position within a radial layer. For clarity, the notation utilized by Yang (1997) is replicated. Suppose that the radial and tangential indices for node i are I and A respectively ($I = 1, \dots, nr + 1$; $A = 1, \dots, 10(mt^2 + 1) + 2$). The piecewise linear shape functions $N_i(\mathbf{r})$ can then be written as:

$$N_i(\mathbf{r}) = M_I(r)L_A(\theta, \phi) \quad (\text{E.37})$$

where $M_I(r)$ and $L_A(\theta, \phi)$ are the radial and tangential shape functions respectively.

The radial shape function is defined as a simple one dimensional linear shape function,

where:

$$M_I(r) = \begin{cases} \frac{r_{I-1} - r}{r_{I-1} - r_I} & \text{when } r_I \leq r < r_{I-1} \\ \frac{r - r_{I+1}}{r_I - r_{I+1}} & \text{when } r_{I+1} \leq r < r_I \\ 0 & \text{otherwise} \end{cases} \quad (\text{E.38})$$

Here, r_I refers to the radius at the I -th radial node. Note that $I = 1$ at the upper boundary.

Linear shape functions are also utilized for the spherical tangential basis function $L_A(\theta, \phi)$, with the value varying linearly, from 1 to 0, as the point (θ, ϕ) moves along a given great circle arc, from node A , towards the outer edges of the surrounding spherical triangles. The function forms the shape of a tent with a pentagonal or hexagonal base centered at node A , depending on whether node A is surrounded by five or six spherical triangles (see Section E.3). $L_A(\theta, \phi) = 0$ outside these triangles.

This relationship is consistent in each and every triangle, in that, as the angular distance from one vertex of the triangle increases along a great circle, the shape function associated with the node decreases linearly. Consequently, local quantities, defined on a triangle, can be utilized, as opposed to a basis function defined globally. The local quantities used for this purpose in TERRA are the barycentric coordinates.

Spherical Barycentric Coordinate System

Spherical barycentric coordinates are denoted by a triplet (η_1, η_2, η_3) , where the subscripts indicate vertices of a spherical triangle (see Figure E.5). A barycentric coordinate, η_α , on a triangle, T , is a measure of the proximity to the vertex α . η_1, η_2 , and η_3 take on the values




Figure E.5: Convention for labeling vertices in (a) a spherical triangle and (b) the sub-triangles from the dyadic refinement. The latter also shows the convention for labeling subtriangles. Modified from Figure 4.5 in Baumgardner (1983).

$(1,0,0)$, $(0,1,0)$, and $(0,0,1)$ respectively at vertices 1, 2, and 3 of T . They vary linearly from 1, at the vertex, to 0, as the point in question moves along a given great circle from a node toward the facing edge. These barycentric coordinates have a value of zero outside of their region of influence, i.e. outside of their 7 point tangential stencil, and 21 point full stencil (both tangential and radial).

The location on any point on the triangle T can be described completely in terms of its barycentric coordinates. They are, therefore, not only the shape functions associated with the three triangle vertices, but they also provide the coordinates of any point in the triangle. This dual nature is simply a consequence of using a linear function. An important concept to grasp is the relationship between the barycentric coordinates $\eta = (\eta_1, \eta_2, \eta_3)$ on a triangle T , and the barycentric coordinates η^1, η^2, η^3 , and η^4 on the four subtriangles, T_1, T_2, T_3 , and T_4 of the dyadic subdivision of T . This can be expressed as:

$$\eta = (\eta_1, \eta_2, \eta_3) = \begin{cases} \frac{1}{2} + \frac{1}{2}\eta_1^1 & \frac{1}{2}\eta_2^1 & \frac{1}{2}\eta_3^1 & \text{on } T_1 \\ \frac{1}{2}\eta_1^2 & \frac{1}{2} + \frac{1}{2}\eta_2^2 & \frac{1}{2}\eta_3^2 & \text{on } T_2 \\ \frac{1}{2}\eta_1^3 & \frac{1}{2}\eta_2^3 & \frac{1}{2} + \frac{1}{2}\eta_3^3 & \text{on } T_3 \\ \frac{1}{2} - \frac{1}{2}\eta_1^4 & \frac{1}{2} - \frac{1}{2}\eta_2^4 & \frac{1}{2} - \frac{1}{2}\eta_3^4 & \text{on } T_4 \end{cases} \quad (\text{E.39})$$

with the labeling convention for vertices and triangles displayed in Figure E.5. This relationship is highly beneficial in establishing relations, called recursion formulas, for various operators between different refinement levels of the tangential grid. They therefore provide an efficient way in evaluating the various operators at different grid levels, which will be highly beneficial in the context of a multigrid.

Having provided an overview of TERRA's finite element basis functions, the natural progression from here would involve a complete description of how the computational operators are expressed in terms of these basis functions. However, this is not an essential prerequisite to the work presented in this thesis and consequently, such a description is not included. The information provided here is sufficient background to the modifications made to the code as a part of this thesis. However, should the reader wish to know more about the derivation of TERRA's computational operators, a detailed discussion can be found in Baumgardner (1983) and Yang (1997). Consequently, rather than simply repeat the mathematics here, the reader is directed towards the original works.

E.6 Solution Strategy for Force-Balance Equation

Having introduced the finite element equations in some detail, the techniques involved in their solution are next described.

Equations (E.34) and (E.35) could be solved through either direct or iterative approaches. The direct methods yield answers in a finite number of operations. Examples include Gauss elimination, Thomas algorithm etc... However, such methods are only suitable if $\mathbf{A}$ is a linear operator. They are also highly impractical when a large number of unknowns are present, since they lead to large matrices which generate high memory demands. Iterative methods on the other hand are suitable for both linear and non-linear equations, in addition to being efficient at dealing with a substantial number of degrees of freedom. Examples include Gauss-Seidel methods, relaxation methods, conjugate gradient methods (CGM), and generalized minimal residual (GMRES) methods. Here, solutions are obtained through a number of iterative steps, with accuracy being increased by an increase in the number of iterations. An iterative procedure of this kind is employed by TERRA to solve the force-balance equation.

E.6.1 A Pressure Correction Approach

One of the most popular methods for solving a Stokes like system, where inertial forces are small compared to viscous forces, is the pressure correction (PC) approach. This is the case with TERRA where a ‘Uzawa’ type PC approach is coupled with a conjugate gradient algorithm to solve Equations (E.34) and (E.35). The basis of this approach is that the velocity and pressure determined by solving Equation (E.34) alone should be corrected until Equation (E.35) is satisfied. The utilized algorithm was originally proposed by Verfurth (1984) and is outlined in detail by both Atanga & Silvester (1992) and Ramage & Wathen

$\mathbf{u}_0 = \mathbf{u}$	
$p = p_0$	
$\mathbf{A}\mathbf{u}_0 - Gp_0 = \mathbf{f}$	Solve for $\mathbf{u}_0$ with initial guess $p = p_0$
$\mathbf{r}_0 = G^T \mathbf{u}_0$	Calculate initial residual
$i = 1$	Initialize iteration count
Do $i = 1, N_1$	Loop until residual is small
If $i = 1$ then	
$\mathbf{s}_1 = \mathbf{r}_0$	Set pressure search direction
Else	
$\delta = \frac{\langle \mathbf{r}_{i-1}, \mathbf{r}_{i-1} \rangle}{\langle \mathbf{r}_{i-2}, \mathbf{r}_{i-2} \rangle}$	
$\mathbf{s}_i = \mathbf{r}_{i-1} + \delta \mathbf{s}_{i-1}$	Set pressure search direction
End if	
$\mathbf{A}\mathbf{v}_i = G\mathbf{s}_i$ for $\mathbf{v}_i$	Solve for velocity search direction
$\alpha = -\frac{\langle \mathbf{r}_{i-1}, \mathbf{r}_{i-1} \rangle}{\langle \mathbf{s}_i, G^T \mathbf{v}_i \rangle}$	Set step length
$p_i = p_{i-1} + \alpha \mathbf{s}_i$	Update pressure
$\mathbf{u}_i = \mathbf{u}_{i-1} + \alpha \mathbf{v}_i$	Update velocity
$\mathbf{r}_i = \mathbf{r}_{i-1} + G^T \mathbf{v}_i$	Update residual
If $(\ \mathbf{r}_i\ < \epsilon_d)$	Exit loop
End Do	
$\mathbf{u} = \mathbf{u}_{N_1}$	
$p = p_{N_1}$	

Table E.1: The pressure correction, conjugate gradient algorithm utilized in ‘TERRA’.

(1994). A brief overview of the algorithm is presented here.

Equation (E.34) implies:

$$G^T \mathbf{u} - G^T \mathbf{A}^{-1} G p = G^T \mathbf{A}^{-1} \mathbf{f} \quad (\text{E.40})$$

It is clear that if a pressure p^* can be found such that:

$$-G^T \mathbf{A}^{-1} G p = G^T \mathbf{A}^{-1} \mathbf{f} \quad (\text{E.41})$$

then $G^T \mathbf{u}^* = 0$ and the continuity equation is satisfied, where $\mathbf{u}^*$ is the related velocity obtained from solving Equation (E.34).

One obvious way to solve this system is to solve first Equation (E.41) for p and then Equation (E.34) for $\mathbf{u}$. This approach requires the assembly and storage of the Schur

complement $(G^T \mathbf{A}^{-1} G)$ as well as $\mathbf{A}$, which is not appealing. Additionally, computing $\mathbf{A}^{-1}$ is an expensive procedure. However, by utilizing a conjugate gradient algorithm, these equations can be solved simultaneously, which means that the Schur complement need not be stored. The *pseudo-code* for this algorithm is presented in Table E.1.

The inner velocity systems in Table E.1 are solved using a highly efficient multigrid scheme, which is described next. Appendix F provides a brief overview of multigrid schemes. The following section simply covers the multigrid formulation employed within TERRA. A reader not familiar with the multigrid concept should, therefore, refer to Appendix F before continuing.

TERRA's Multigrid Formulation

For simplicity, TERRA's multigrid algorithm is presented for a two level system, consisting of a fine grid (grid spacing h) and one coarse grid (grid spacing $2h$). Suppose we are given an estimate, $\mathbf{u}_0$, for the solution of E.36. The residual is then given by:

$$\mathbf{r}_h = \mathbf{f}_h - \mathbf{A}_h \mathbf{u}_0 \quad (\text{E.42})$$

The solution to the equation:

$$\mathbf{A}_h \mathbf{e}_h = \mathbf{r}_h \quad (\text{E.43})$$

is sought, where $\mathbf{e}_h$ represents the error associated with the approximate solution - the correction that should be added to $\mathbf{u}_0$ to improve the estimate for $\mathbf{u}$. A crucial aspect of the multigrid method is to find a coarser grid correction field $\mathbf{e}_{2h}$ from:

$$\mathbf{A}_{2h} \mathbf{e}_{2h} = \mathbf{r}_{2h} \quad (\text{E.44})$$

where $\mathbf{A}_{2h}$ and $\mathbf{r}_{2h}$ are the forward operator and the residual vector on the coarser grid. As noted in Appendix F, a restriction operator is required to generate the coarser grid

representation, $\mathbf{r}_{2h}$, of the residual from the fine grid residual, $\mathbf{r}_h$. Construction of the coarse grid forward operator, $\mathbf{A}_{2h}$ is another important issue. Both issues are discussed in the following sections, with TERRA utilizing an inter-grid transfer method known as *matrix dependent transfer*.

Once we have $\mathbf{A}_{2h}$ and $\mathbf{r}_{2h}$, we can determine $\mathbf{e}_{2h}$ either through a relaxation method, or, when the matrix is manageable, solving the equation exactly. A *pseudo-inverse* method is utilized in TERRA at this stage (see Baumgardner 1983, for further details). The process of finding $\mathbf{e}_{2h}$ is called the coarse grid correction.

The fine grid representation $\mathbf{e}'_h$ of $\mathbf{e}_{2h}$ can be obtained by utilizing an interpolation operator. When compared to the correction scheme described in Appendix F, TERRA incorporates minor subtleties at this point. To improve the correction field, $\mathbf{e}'_h$, thus obtained a second order residual is formed on the fine-grid:

$$\mathbf{r}'_h = \mathbf{r}_h - \mathbf{A}_h \mathbf{e}'_h \quad (\text{E.45})$$

A second order correction term $\mathbf{e}''_h$ is then found by relaxing (using line-Jacobi relaxation) on the equation:

$$\mathbf{A}_h \mathbf{e}''_h = \mathbf{r}'_h \quad (\text{E.46})$$

The correction field on the fine grid is then given by:

$$\mathbf{e}_h = \mathbf{e}'_h + \mathbf{e}''_h \quad (\text{E.47})$$

One multigrid iteration is completed when this term is added to $\mathbf{u}_0$. The corrected solution is then fed back into Equation (E.42) and the iterating process is continued until the residual becomes acceptably small. This is an example of the *V-cycle* illustrated in Figure F.2 of Appendix F.

E.6.2 Matrix Dependent Transfer

Alcouffe et al. (1981) note that the usual interpolation schemes employed for intergrid transfers in the multigrid method (the restriction and interpolation operators) yield poor convergence for solving problems with strongly discontinuous coefficients in the governing differential equations. Based on this insight, a more appropriate intergrid transfer scheme, called ‘matrix dependent transfer’, was developed. Various studies (e.g. Wagner 1994) have since demonstrated that the idea of matrix dependent transfer is frequently the essential ingredient of a successful formulation to yield convergent solutions in strongly discontinuous or variable coefficient problems.

Since the viscosity varies by orders of magnitude across Earth’s mantle, Yang (1997) exploited this idea within TERRA’s multigrid formulation and adapted the original code to include matrix dependent intergrid transfers. The central idea in this new approach is that the interpolation, $\mathbf{P}$, and restriction, $\mathbf{R}$, operations, as well as the coarse grid operators themselves, be obtained from the fine grid operator, $\mathbf{A}$. The concept is explained briefly in Yang (1997), with a more detailed account presented in Yang & Baumgardner (2000). It is important to note, that in this formulation, the restriction operator, $\mathbf{R}$, is derived from the transpose of the interpolation operator, $\mathbf{P}$, the so-called canonical choice:

$$\mathbf{R} = \mathbf{P}^T \tag{E.48}$$

The method has proven to be a crucial element in the successful application of the multigrid method to variable viscosity mantle convection simulations with TERRA.

E.6.3 Coarse Grid Operators

The coarse grid forward operator, $\bar{\mathbf{A}}$, can be derived through either:

1. Discretization coarse grid approximation - the coarse grid forward operator is assembled from the elementary radial and tangential operators, derived from the finite element shape functions on the coarse grid (as occurs on the finest grid), or,
2. Galerkin coarse grid approximation - where the prescription:

$$\bar{\mathbf{A}} = \mathbf{R}\mathbf{A}\mathbf{P} \quad (\text{E.49})$$

is utilized. Since $\mathbf{R}$ is customarily chosen as $\mathbf{P}^T$, the coarse grid operator, formed via the Galerkin approximation, has the same definiteness as the fine grid operator.

The chief problem with the first method is that the viscosity field is often difficult to approximate on the coarse grid levels. The Galerkin scheme on the other hand naturally avoids this issue, since it does not require any coarse grid viscosity field. The coarse grid forward operator is obtained solely from the fine grid operator, prolongation matrix and restriction matrix (utilizing the matrix dependent transfer scheme described above). Also, as noted by Wesseling (1992), it is more appropriate on coarse grids not to sample variable coefficients only on a sparse set of points, which is what the discretization coarse grid approximation does, but to take suitable averages, which is what the Galerkin coarse grid approximation does automatically. Consequently, the Galerkin Coarse Grid scheme, in conjunction with the matrix dependent intergrid transfer method, is utilized to generate the coarse grid forward operators in TERRA.

E.6.4 Relaxation Method

Line Jacobi relaxation is employed within TERRA, as part of the multigrid scheme, to solve Equation (E.36), using the unmodified forward operator, at each grid level. This relaxation is performed in the radial direction, since it is common for material properties to show a stronger variation radially than tangentially within the spherical shell. Typically, one to four sweeps of the line Jacobi procedure are made with a weighting factor of 2/3. Various other

relaxation methods have been tested by Yang (1997), including Gauss-Seidel relaxation, however, superior convergence rates were yielded by the Jacobi relaxation algorithm.

E.7 The Energy Equation and Time Integration

The methods used to treat the advection, thermal conduction and time integration of the energy conservation equation (E.3) are described in this section, in terms of the discretization presented in the previous sections.

E.7.1 Cell-Wall Advection

The advection of heat is accomplished through a flux-form scheme known as ‘cell-wall advection’. Fluxes are calculated at element walls, with advection only occurring between elements that share a common face. The direction of movement is determined by the sign of the normal velocity, with the rate of movement proportional to the magnitude of normal velocity integrated over the area of the ‘common-face’. The density, temperature and specific heat of the cell, for which the integrated normal velocity is outward, are also essential in determining the rate of advection.

As may be noted from Figure E.2, 3-D icosahedral elements take the form of triangular prisms with three planar faces and two spherical ends. For the remainder of this section, advection through the planar faces will be termed ‘tangential’, while advection through spherical ends is termed ‘radial’. In both cases, the integral of the normal velocity over the face area is determined by multiplying the normal velocities at the vertices of the face by appropriate weighting factors and summing. For planar faces, the weighting factors for nodes at the inner radius r_i of the face are given by:

$$\frac{(2r_i + r_0)(r_0 - r_i)\phi}{12} \quad (\text{E.50})$$

while the weights for nodes at the outer radius r_0 are:

$$\frac{(r_i + 2r_0)(r_0 - r_i)\phi}{12} \quad (\text{E.51})$$

where ϕ is the angle subtended by the face. For spherical faces, the weighting factors are given by:

$$A \left(\frac{x_{ki} + \frac{1}{2}x_{ki'} + \frac{1}{2}x_{ki''}}{2} \right) \quad (\text{E.52})$$

Here, A represents the area of the spherical triangle at unit radius, while x_{ki} denotes the Cartesian coordinate k of triangle vertex i . i' and i'' denote the triangle vertices different from i .

The advection calculation itself involves the following stages:

1. Resolve the velocity field into radial and tangential components at each node.
2. Compute integrated normal velocities for each face.
3. Loop over all elements, from the surface to the core, calculating the rate of change of mass and heat, by accumulating contributions across all faces. The rate of mass advection through a given face is the integrated normal velocity times the density of the donor cell (the cell for which the integrated normal velocity is outward). Similarly, the rate of heat advection is the product of the donor cell density, specific heat and temperature, and the face integrated normal velocity.

The proposed method does have its drawbacks in that it introduces a slight error in integration of the normal velocity on spherical faces. However this error is proportional to the element size, and consequently, becomes negligible as the discretization becomes fine, when spherical triangles show only a small variation in normal direction. Clearly therefore, the advantages of the method, in being simple to implement and highly efficient, outweigh the disadvantages.

E.7.2 Thermal Conduction

The modeling of thermal conduction is achieved through a method that is analogous to that utilized in tracking thermal advection. It is based around element-to-element conductivity factors, which are determined from the geometry of the grid. The rate of conductive heat flow between adjacent elements is found by multiplying the temperature difference between the elements by an appropriate conductivity factor.

In an manner similar to cell-wall advection, tangential (planar faces) and radial conduction (spherical faces) is treated separately. The conductivity factor, ξ_r , for radial conduction is given by:

$$\xi_r = \frac{kA\bar{r}_1\bar{r}_2}{(\bar{r}_2 - \bar{r}_1)} \quad (\text{E.53})$$

where k is the thermal conductivity and A is the area of the spherical face at unit radius. $\bar{r}_1$ and $\bar{r}_2$ are defined by:

$$\bar{r}_1 = \frac{2(r_{i-1}^2 + r_{i-1}r_i + r_i^2)}{3(r_{i-1} + r_i)} \quad (\text{E.54})$$

and:

$$\bar{r}_2 = \frac{2(r_i^2 + r_ir_{i+1} + r_{i+1}^2)}{3(r_i + r_{i+1})} \quad (\text{E.55})$$

In the above equations, r_{i-1} represents the inner radius of the lower element, r_i the radius of the common face, and r_{i+1} the outer radius of the upper element.

Temperature boundary conditions are handled in the same fashion as element temperatures - the value for $\bar{r}_2$ at the outer shell boundary is simply the outer radius of the shell, while the value for $\bar{r}_1$ at the inner shell boundary is the inner radius of the shell. It has been verified (Baumgardner, 1983) that this discrete formulation yields a result which agrees exactly with the analytical solution.

For tangential conduction, the conductivity factor ξ_η is given by:

$$\xi_\eta = \frac{k\phi(r_0 - r_i)}{\delta} \quad (\text{E.56})$$

where ϕ is the half-angle subtended by the face with respect to the center of the sphere, r_0 the outer radius of the face, r_i the inner radius of the face and δ represents the distance between cell centers when both are projected to unit radius. The advocated method yields results that are within $< 0.5\%$ of the analytical solution (Baumgardner, 1983). Since tangential conduction only plays a minor role in the mantle convection problem, the accuracy indicated here is deemed satisfactory.

E.7.3 Time Integration

The discrete time integration procedure implemented provides second-order accuracy, but requires derivatives at two points in each time interval. The scheme is a member of the Runge-Kutta family, which can be recast into an m -stage scheme of the form:

$$\Delta T^{n+1} = \alpha_i \Delta t r (T^n + \Delta T^{n+i-1}) \quad i = 1, m \quad (\text{E.57})$$

where T is the temperature, r represents the residual in the energy equation and Δt is the time step. The coefficients α are chosen according to desired properties, such as the temporal order of accuracy. The coefficients implemented in TERRA's second-order scheme are $\alpha_1 = 0.5$ and $\alpha_2 = 1$. Written out in full, the time-stepping algorithm is as follows:

$$\begin{aligned} \Delta T^n &= 0 \\ \Delta T^{n+1} &= \alpha_1 \Delta t r (T^n + \Delta T^n) \\ \Delta T^{n+2} &= \alpha_2 \Delta t r (T^n + \Delta T^{n+1}) \end{aligned} \quad (\text{E.58})$$

or alternatively:

$$T_0 = \text{known} = T^n$$

$$T_1 = T_0 + \alpha_1 \Delta t r(T_0) \quad (\text{E.59})$$

$$T_2 = T_0 + \alpha_2 \Delta t r(T_1) = T^{n+1}$$

In the above equations, r is the residual of the energy equation. The time step is continually adjusted in the overall calculation, with the intention that only one iteration of the multigrid algorithm is needed to maintain the residual error in the force-balance equation below a specified level. At the same time, an upper bound is set on the time step such that the mass advected through any cell in the grid in a single time step never exceeds the volume of the cell.

E.8 Parallelization

E.8.1 Overview

Terra was parallelized by Bunge & Baumgardner (1995) using explicit message passing software known as Parallel Virtual Machine (PVM). The parallelization has since been modified (Yang 1997) with the more advanced message passing software, Message Passing Interface (MPI), now being utilized. The geometry and scale of the convective problems under study mean that obtaining reasonably resolved results whilst running the code on a single processor is simply out of the question. The code's parallelization has opened the door to much more realistic simulations, at 'Earth' like Rayleigh numbers. Indeed, of all the modifications made to TERRA, this is possibly the most significant.

E.8.2 Technicalities

Parallelization is accomplished through a procedure termed '*domain decomposition*'. The existing computational domain - the icosahedral grid - is *decomposed* into smaller *subdomains* and spread across a number of processes.




Figure E.6: Subdomain processor mapping in TERRA - $mt/nt = 2$. In (a) $nd = 10$, while in (b) $nd = 5$. The diamonds have been projected on to a flat surface and solid black lines define their boundaries, while dashed lines represent subdomain boundaries. The number within each subdomain denotes the MPI rank of the process to which the subdomain is mapped.

The first step is to divide each diamond into a series of subdiamonds/subdomains. As stated in section E.3.1, the number of grid intervals along an icosahedral diamond edge is referred to as mt . A second parameter is used to define the size of the subdomains: nt - the number of grid intervals along the edge of a local subdomain. Subdomains, like diamonds have $i1$ and $i2$ axes that have the same length. The values of mt and nt must be such that a diamond is divided into 1, 4, 16, 32 or subdomains, i.e. mt must be a power of 2 and nt must also be a power of two less than or equal to mt .

The next step in the decomposition is selecting the number of subdomains to distribute to each process. This is defined by the parameter nd - the number of diamonds from which subdomains will be mapped onto the processors. nd can have a value of 5 or 10. If $nd = 5$ only northern hemisphere diamonds are mapped onto the first half of the processes and southern hemisphere diamonds to the second half. If $nd = 10$ each process owns one subdomain from each of the ten diamonds. This is illustrated in Figure E.6.

Although conceptually simple, such domain decomposition unfortunately has one minor disadvantage - subdomains extend throughout the radial dimension. This generates long, thin subdomains with a large surface area. Such a large surface area is undesirable since it leads to excessive message passing that restricts the performance and ‘speedup’ of the code. Nonetheless, as stated above, parallelization has allowed us to study problems at resolutions that would otherwise be unattainable, and, consequently this minor drawback can be overlooked.

E.9 Visualization

TERRA produces extremely large datasets. Consequently, visualization is a difficult task - facilities and programs must be able to cope with extremely large amounts of data, in addition to being able to display the full 3-D output, i.e. spherical surfaces at various depths.

Several packages have been designed to deal with these issues and consequently results can be displayed in a variety of forms. The images presented in this thesis have been generated by an ‘in-house’ program known as ‘Mantlevis’. Mantlevis was written by Dr. Andy Heath at Liverpool in the late 1990’s. It runs in a client-server manner, with the server being the output graphic device, which interprets the OpenGL graphic commands, while the clients are the processors that produce the OpenGL graphic commands by undertaking the rendering. It runs in parallel, over the same processors as a simulation in TERRA. This has significant advantages in that large datasets can be viewed quickly and efficiently. Additionally, the data does not need to be migrated away from the computational cluster, whilst perhaps more importantly, a cluster capable of running TERRA is capable of running mantlevis. This negates the need for extra visualization rendering hardware.

Appendix F

Multigrid

F.1 Introduction

In numerical modeling one often needs to solve a linear equation of the form:

$$\mathbf{A}\mathbf{u} = \mathbf{f} \tag{F.1}$$

for $\mathbf{u}$, where the rank of the matrix $\mathbf{A}$ is extremely large. Since direct matrix inversion is out of the question for such large systems, an iterative approach is commonly employed. The multigrid method is one such example, combining various techniques and concepts to generate one of the most powerful numerical algorithms available today. The fundamental elements of the method are covered here, although more complete descriptions can be found in Brandt (1984) and Briggs et al. (2000). Note that the boldface notation usually employed to denote vectors and tensors is dropped for the remainder of this appendix.

F.2 The Residual Equation

Suppose the system $Au = f$ has a unique solution and that v is a computed approximation to u . There are two important measures of v as an approximation to u :

1. *The error*, given by $e = u - v$.
2. *The residual*, given by $r = f - Av$.

Unfortunately, the error is just as inaccessible as the exact solution itself. The residual, however, is a computable measure of how well v approximates u . It represents the amount by which the approximation, v , fails to satisfy the original problem, $Au = f$. Using the definitions for r and e and noting that $Au = f$, an extremely important relationship can be derived between the error and the residual:

$$Ae = r \tag{F.2}$$

This is termed the *residual equation*. It states that the error satisfies the same set of equations as the unknown u , when f is replaced by the residual, r . It can therefore be used to great advantage. Suppose an approximation, v , has been computed by some method. The residual can be easily computed through $r = f - Av$. v can then be improved by solving the residual equation for e and subsequently computing a new approximation using the definition of the error:

$$u = v + e \tag{F.3}$$

In practice, this method must be applied more carefully than has been indicated here. Nonetheless, this idea of residual correction is highly important in all that follows.

F.3 Relaxation Methods

We now turn our attention to relaxation methods, more commonly known as iterative methods. It is a well known fact that relaxation schemes work very well at reducing the solution error for the first few iterations. Inevitably however, convergence slows and the entire scheme appears to stall. A close inspection of this behavior reveals that the convergence

rate is a function of the error field frequency, i.e. the gradient of the error from node to node. The rapid decrease in error during the early iterations is due the efficient elimination of the oscillatory modes of that error. However, once the oscillatory modes have been removed, relaxation is far less effective at reducing the remaining smooth components.

Many relaxation schemes possess this so-called *smoothing property* and it is a serious limitation. Fortunately however, this shortcoming can be overcome and the remedy is one of the pathways to multigrid. How therefore can these methods be modified to make them effective on all error components?

One way to improve a relaxation scheme, at least in its early stages, is to use a good initial guess. A well-known technique for obtaining an improved initial guess is to perform some preliminary iterations on a coarse grid. Relaxation on a coarse grid is less expensive, since fewer unknowns need to be updated. This line of reasoning suggests that coarse grids might be worth considering.

With the coarse grid idea in mind, we can think more carefully about its implications. Recall that most basic relaxation schemes suffer in the presence of smooth error components. Assume that a particular scheme has been applied until only smooth error components remain. We now ask what these smooth components will look like on a coarser grid. The answer is illustrated in Figure F.1, where a smooth wave has been projected from a fine grid with $n = 16$ points, directly on to a coarser grid with $n = 8$ points. We see that a smooth wave appears more oscillatory on a coarse grid. In other words, in passing from the fine grid to the coarse grid, a mode becomes more oscillatory. This suggests that when relaxation begins to stall, signalling the predominance of smooth error modes, it is advisable to move to a coarser grid; there, the smooth error modes appear more oscillatory and relaxation will be more effective. The question is: how do we move to a coarser grid and relax on the more




Figure F.1: A wave displayed on both fine (upper) and coarse (lower) one-dimensional grids. The coarse grid ‘sees’ a wave that is more oscillatory than it appears on the fine grid

oscillatory error modes?

It is at this point that the multigrid technique begins to come together. The two main points can be summarized as follows:

1. The smoothing property of many relaxation schemes leads us to consider using coarser grids during the computation to focus the relaxation on the oscillatory components of the error.
2. In addition, there seems to be good reason to involve the residual equation in some way, since it allows us to relax directly on the error. There is also another argument that justifies the use of the residual equation: relaxation on the original equation $Au = f$ with an arbitrary initial guess v , is equivalent to relaxing on the residual equation $Ae = r$ with the specific initial guess $e = 0$. Such an intimate connection further motivates the use of the residual equation.

F.4 Coarse Grid Correction

The previous section highlights the fact that relaxation sweeps rapidly reduce high frequency error components. Smoother components should then be reduced by coarse-grid

approximation and relaxation, on a grid with mesh-size $H = 2h$, for example, where h is the characteristic mesh-size of the fine grid. Generally, for any linear fine-grid equation $A^h u^h = f^h$, and any approximate solution v^h , the error, $e^h = u^h - v^h$ satisfies:

$$A^h e^h = r^h \quad \text{where} \quad r^h = f^h - A^h v^h \quad (\text{F.4})$$

This error, e^h , can therefore be approximated by the coarse-grid function e^H , which satisfies:

$$A^H e^H = I_h^H r^h \quad (\text{F.5})$$

where A^H is some coarse-grid approximation to A^h , and I_h^H is a fine-to-coarse transfer operator, more commonly known as a *restriction* operator. That is, $I_h^H r^h$ is a coarse-grid function whose value at each point is a certain weighted average of values of r^h at neighboring fine-grid points.

Having obtained an approximate solution e^H to Equation (F.5), we use it as a correction to the fine-grid solution. Namely, we replace:

$$v^h \leftarrow v^h + I_H^h e^H \quad (\text{F.6})$$

where I_H^h is a coarse-to-fine *interpolation* or *prolongation* operator. That is, at each fine-grid point, the value of $I_H^h e^H$, designed to approximate the error e^h , is interpolated from values of e^H at neighboring coarse-grid points. Linear interpolation can be used in most cases. An important point to note here is that interpolation is most effective when the error is smooth - if the error is oscillatory, even a very good coarse-grid approximation may produce an interpolant that is not very accurate. Fortunately, this provides a complement to relaxation, which is most effective when the error is oscillatory.

The whole process of calculating $I_H^h r^h$, solving Equation (F.5) and interpolating the correction (F.6) is called a *coarse grid correction*. This procedure is the basis of the so-called




Figure F.2: Schedule of grids for various multigrid cycles, all on four levels.

correction scheme (CS). Having relaxed on the fine grid until convergence deteriorates, we relax on the residual equation on a coarser grid to obtain an approximation to the error itself. We then return to the fine grid to correct the original approximation. It is important to appreciate how well the various elements complement each other. Relaxation on the fine grid eliminates the oscillatory components of the error, leaving a relatively smooth error that can be well-approximated on a coarser grid, where solution is much cheaper. Assuming the residual equation can be solved accurately on a coarse grid, it is still important to transfer the error accurately back to the fine grid. Because the error is smooth, interpolation will work very well and the correction to the fine-grid solution should be effective.

The two grid scheme, as outline above, leaves one looming procedural question: what is the best way to solve the coarse grid problem $A^H e^H = r^H$? An approximate solution to the coarse-grid equation (F.5) can be obtained by employing the above solution process recursively; i.e. Equation (F.5) is itself solved by relaxation sweeps combined with a still coarser grid correction. We thus have a sequence of $l > 1$ grids with grid spacings $h, 2h, 4h, \dots, Lh = 2^{l-1}h$. Schemes of this nature can be implemented in numerous forms, the most common of which are illustrated in Figure F.2. A brief run through of the so-called

V -cycle is provided here. The other cycles are extremely similar in concept, excluding the scheduling of grids, therefore, to avoid repetition, only the V -cycle is described fully.

F.4.1 V -Cycle Scheme

The V -cycle multigrid scheme takes on the following form:

- Relax v_1 times on $A^h v^h = f^h$ with initial guess v^h .
- Compute fine-grid residual, $r^h = f^h - A^h v^h$.
- Restrict residual to coarse grid by $r^{2h} = I_h^{2h} r^h$.
 - Relax v_1 times on $A^{2h} e^{2h} = r^{2h}$ with initial guess $e^{2h} = 0$.
 - Compute $r^{4h} = I_{2h}^{4h} r^{2h}$.
 - * Relax v_1 times on $A^{4h} e^{4h} = r^{4h}$ with initial guess $e^{4h} = 0$.
 - * Compute $r^{8h} = I_{4h}^{8h} r^{4h}$.
 - Solve $A^{Lh} e^{Lh} = r^{Lh}$ on Ω^{Lh}
 - * Correct $e^{4h} \leftarrow e^{4h} + I_{8h}^{4h} e^{8h}$.
 - * Relax v_2 times on $A^{4h} e^{4h} = r^{4h}$ with initial guess e^{4h} .
 - Correct $e^{2h} \leftarrow e^{2h} + I_{4h}^{2h} e^{4h}$.
 - Relax v_2 times on $A^{2h} e^{2h} = r^{2h}$ with initial guess e^{2h} .
- Correct $v^h \leftarrow v^h + I_{2h}^h e^{2h}$.
- Relax v_2 times on $A^h v^h = f^h$ with initial guess v^h .

The algorithm telescopes down to the coarsest grid, then works its way back to the finest. Figure F.2 shows the schedule for the grids in the order in which they are visited. Because of the pattern displayed in this diagram, the algorithm is called the V -cycle.

F.5 The Full Approximation Storage Scheme

The Full Approximation Scheme (or Full Approximation Storage - FAS) is a widely used version of multigrid inter-grid transfers. It has been mainly used in solving non-linear

problems, but it has so many other applications that it should perhaps be used in most advanced programs.

In the FAS Scheme, exactly the same steps are performed as the Correction Scheme, but in terms of another coarse-grid variable. Instead of e^H , we use:

$$u^H = I_h^H v^h + e^H \quad (\text{F.7})$$

as the coarse-grid unknown function. This coarse grid variable, u^H approximates $I_h^H u^h$, the full intended solution represented on the coarse grid, hence the name ‘Full Approximation Scheme’. The FAS coarse-grid equations, derived from Equations (F.5) and (F.7), are:

$$A^H u^H = \hat{f}^H \quad (\text{F.8})$$

where:

$$\hat{f}^H = A^H(I_h^H v^h) + I_h^H r^h \quad (\text{F.9})$$

Having obtained an approximate solution, $\tilde{u}^H$, to Equation (F.8), the approximate coarse-grid correction is of course:

$$e^H = \tilde{u}^H - I_h^H v^h \quad (\text{F.10})$$

hence, the FAS interpolation back to the fine grid, equivalent to Equation (F.6), is:

$$v^h \leftarrow v^h + I_H^h(\tilde{u}^H - I_h^H v^h) \quad (\text{F.11})$$

To use directly:

$$v^h \leftarrow I_H^h \tilde{u}^H \quad (\text{F.12})$$

would be worse, of course, since it would introduce the interpolation errors of the full solution u^H , instead of the interpolation errors of only the correction e^H .

F.6 Conclusions

In summary, the multigrid method is a remarkable synthesis of ideas and techniques that individually have been well known and used for a long time. Taken alone, many of these ideas have significant defects. The multigrid however, is a technique for integrating them so that they can work together in a way that overcomes these limitations. The result is a very powerful tool for numerical analysis.

Although not all elements of the technique have been covered here, the basic concepts have been introduced. The aim of this appendix is not to provide an in depth analysis of the method as a whole. It is merely intended to provide the reader with some background, which will prove useful in understanding this thesis. The more advanced features of the technique, specific to the work contained herein, are included amongst various chapters, wherever appropriate.

