

©Copyright 2013
Grady Lemoine

Numerical Modeling of Poroelastic-Fluid Systems Using High-Resolution Finite Volume Methods

Grady Lemoine

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2013

Reading Committee:

Randall J. LeVeque, Chair

Ulrich L. Hetmaniuk

M. Yvonne Ou

Program Authorized to Offer Degree:
Applied Mathematics

University of Washington

Abstract

Numerical Modeling of Poroelastic-Fluid Systems Using High-Resolution Finite Volume Methods

Grady Lemoine

Chair of the Supervisory Committee:
Professor Randall J. LeVeque
Applied Mathematics

Poroelasticity theory models the mechanics of porous, fluid-saturated, deformable solids. It was originally developed by Maurice Biot to model geophysical problems, such as seismic waves in oil reservoirs, but has also been applied to modeling living bone and other porous media. Poroelastic media often interact with fluids, such as in ocean bottom acoustics or propagation of waves from soft tissue into bone.

This thesis describes the development and testing of high-resolution finite volume numerical methods, and simulation codes implementing these methods, for modeling systems of poroelastic media and fluids in two and three dimensions. These methods operate on both rectilinear grids and logically rectangular mapped grids. To allow the use of these methods, Biot's equations of poroelasticity are formulated as a first-order hyperbolic system with a source term; this source term is incorporated using operator splitting. Some modifications are required to the classical high-resolution finite volume method. Obtaining correct solutions at interfaces between poroelastic media and fluids requires a novel transverse propagation scheme and the removal of the classical second-order correction term at the interface, and in three dimensions a new wave limiting algorithm is also needed to correctly limit shear waves.

The accuracy and convergence rates of the methods of this thesis are examined for a variety of analytical solutions, including simple plane waves, reflection and transmission of waves at an interface between different media, and scattering of acoustic waves by a poroelastic cylinder. Solutions are also computed for a variety of test problems from the computational poroelasticity literature, as well as some original test problems designed to mimic possible applications for the simulation code.

TABLE OF CONTENTS

	Page
List of Figures	iv
List of Tables	vi
Chapter 1: Introduction	1
1.1 Motivation	1
1.2 Previous work in computational poroelasticity	2
1.3 Contribution of this work	3
1.4 Outline	4
Chapter 2: Governing equations	6
2.1 Two dimensional orthotropic, transversely isotropic poroelasticity	6
2.2 Three dimensional orthotropic poroelasticity	17
2.3 Linear acoustics	22
2.4 Interface conditions	24
2.5 Decay of poroelastic waves in practical problems	26
Chapter 3: Methods for poroelasticity on two-dimensional rectilinear grids	30
3.1 Wave propagation	30
3.2 Operator splitting	32
3.3 Stiff regime and subcharacteristic condition	32
3.4 Numerical software	35
Chapter 4: Results for poroelasticity on two-dimensional rectilinear grids	36
4.1 Analytical plane wave solution	36
4.2 Plane wave convergence results — inviscid	37
4.3 Plane wave convergence results — viscous	41
4.4 Single-material point source results	46
4.5 Heterogeneous domain results	47

Chapter 5:	Methods for poroelastic-fluid systems on two-dimensional mapped grids	52
5.1	Mapped grids	52
5.2	Riemann problems on mapped grids	53
5.3	Modifications to the high-resolution finite volume scheme	59
5.4	Source term and numerical software	63
Chapter 6:	Results for poroelastic-fluid systems on two-dimensional rectilinear grids	65
6.1	Effect of omitting the second-order correction term	66
6.2	Reflected and transmitted waves at a fluid-poroelastic interface	67
6.3	Reflected and transmitted waves at an interface between two poroelastic materials	70
Chapter 7:	Results for poroelastic-fluid systems on two-dimensional mapped grids	73
7.1	Cylindrical scatterer	73
7.2	Femur model	77
Chapter 8:	Methods for poroelastic-fluid systems on three-dimensional mapped grids	81
8.1	Mapped grids in three dimensions	81
8.2	Euler angles	83
8.3	Riemann problems on three-dimensional mapped grids	83
8.4	Entropy function	87
8.5	Shear waves and revised limiter algorithm	88
8.6	Dimensional splitting	91
8.7	Numerical software	93
Chapter 9:	Results for poroelastic-fluid systems on three-dimensional mapped grids	94
9.1	Analytic plane wave solution	94
9.2	Plane wave convergence studies	95
9.3	Demonstration problem	101
Chapter 10:	Conclusions and future work	116
10.1	Summary of results	116
10.2	Opportunities for future work	119
Bibliography	121
Appendix A:	Material properties used in this work	127

Appendix B:	Analytic solution procedure for a time-harmonic plane wave train striking a flat interface	130
Appendix C:	Time-harmonic analytic solution for an acoustic plane wave scattering off an isotropic poroelastic cylinder	133
C.1	Preliminaries	133
C.2	Derivation of a second-order system	134
C.3	Further reduction of unknowns	135
C.4	Application of the Helmholtz decomposition to Biot's equations	137
C.5	Solution of the Helmholtz equations for the wave potentials	139
C.6	Implementation of the solution	147
C.7	Behavior of the series solution	151
Appendix D:	Gallery of energy density plots for cylindrical scatterer cases	158

LIST OF FIGURES

Figure Number	Page
2.1 Fast P and S wave decay lengths for stone media	27
2.2 Slow P wave decay lengths for stone media	27
2.3 Fast P and S wave decay lengths for bone	28
2.4 Slow P wave decay lengths for bone	28
4.1 Sample plane wave initial condition	38
4.2 Comparison of error with and without a limiter	40
4.3 Comparison of wave speeds for full and reduced systems	41
4.4 Operator splitting test results	43
4.5 Results for point source test cases	47
4.6 Vertical velocity and AMR grids for heterogeneous domain test case	48
4.7 Time-histories of the x -direction solid velocity for heterogeneous domain test case	49
4.8 Time-histories of the z -direction solid velocity for heterogeneous domain test case	50
5.1 Example of a mapped grid	53
5.2 Dependence of conditioning of Riemann solve on free parameter ζ	58
5.3 Comparison of classical and fluctuation-based transverse cell updates	61
6.1 Effect of omitting the second-order correction term	67
6.2 Sketches of reflection/transmission test problems	68
7.1 Frequency response of cylindrical scatterer	74
7.2 Energy density for analytical and numerical solutions for cylindrical scatterer	75
7.3 Grid mappings for cylindrical scatterer and femur models	76
7.4 Convergence of cylindrical scatterer test cases	78
7.5 Results of femur simulation	80
8.1 Illustration of yaw, pitch, and roll angles	84
9.1 Brine-sandstone interface surface for demonstration problem	102
9.2 Grid mapping for demonstration problem	103
9.3 Energy density for demonstration problem (inviscid)	106

9.4	Vertical normal stress for demonstration problem (inviscid)	107
9.5	Fluid flow speed for demonstration problem (inviscid)	108
9.6	Vertical normal stress for demonstration problem (inviscid)	109
9.7	Vertical velocity for demonstration problem (inviscid)	110
9.8	Energy density for demonstration problem (viscous)	111
9.9	Vertical normal stress for demonstration problem (viscous)	112
9.10	Vertical normal stress for demonstration problem (viscous)	113
9.11	Vertical velocity for demonstration problem (viscous)	114
D.1	Cylindrical scatterer case 1	158
D.2	Cylindrical scatterer case 2	159
D.3	Cylindrical scatterer case 3	159
D.4	Cylindrical scatterer case 4	160
D.5	Cylindrical scatterer case 5	160
D.6	Cylindrical scatterer case 6	161
D.7	Cylindrical scatterer case 7	161
D.8	Cylindrical scatterer case 8	162
D.9	Cylindrical scatterer case 9	162
D.10	Cylindrical scatterer case 10	163
D.11	Cylindrical scatterer case 11	163
D.12	Cylindrical scatterer case 12	164
D.13	Cylindrical scatterer case 13	164
D.14	Cylindrical scatterer case 14	165
D.15	Cylindrical scatterer case 15	165
D.16	Cylindrical scatterer case 16	166
D.17	Cylindrical scatterer case 17	166
D.18	Cylindrical scatterer case 18	167

LIST OF TABLES

Table Number	Page
4.1 Summary of convergence results for inviscid test cases	39
4.2 Summary of convergence results for viscous, high-frequency test cases	44
4.3 Summary of convergence results for viscous, low-frequency test cases	45
6.1 Convergence results with second-order correction omitted along a line	66
6.2 Convergence results for an acoustic wave impinging on sandstone, with viscosity omitted.	69
6.3 Convergence results for a 10 kHz acoustic wave impinging on sandstone, with viscosity included.	69
6.4 Convergence results for poroelastic waves in shale impinging on sandstone, with viscosity omitted.	71
6.5 Convergence results for 20 kHz poroelastic waves in shale impinging on sandstone, with viscosity included.	72
7.1 Cylindrical scatterer cases chosen for convergence analysis	76
8.1 Effect of using the energy inner product limiter	90
8.2 Summary of convergence results for dimensional splitting	92
9.1 Three-dimensional test cases	96
9.2 Three-dimensional plane wave convergence results, 1-norm, cases 0-19	97
9.3 Three-dimensional plane wave convergence results, 1-norm, cases 20-35	98
9.4 Three-dimensional plane wave convergence results, max-norm, cases 0-19	99
9.5 Three-dimensional plane wave convergence results, max-norm, cases 20-35	100
9.6 Parameters for demonstration problem	101
A.1 Base properties of selected poroelastic media	128
A.2 Derived properties of selected poroelastic media	129

ACKNOWLEDGMENTS

I would like to express my thanks to my advisor, Professor Randall J. LeVeque, for his support and patience regarding this work, as well as for originating the CLAWPACK software framework in which I have implemented all of my numerical work for this thesis. I would also like to thank Professor M. Yvonne Ou of the University of Delaware, for introducing me to poroelasticity, for her invaluable aid on some of the analytical aspects of this project, and for her exemplary hospitality. In addition, I am grateful to the other members of my committee, Professor Pedro Arduino, Professor Ulrich L. Hetmaniuk, and Doctor Thomas J. Matula, for providing insight, context, and constructive criticism. Beyond my committee members, Professor Donna Calhoun of Boise State University also provided valuable insight into formulating finite volume methods on three-dimensional mapped grids, which saved me substantial time and confusion.

Finally, I am grateful to the National Science Foundation and the National Institutes of Health for their support during my doctoral program. This work was supported in part by NIH grant 5R01AR53652-2, and NSF grants DMS-0914942 and DMS-1216732; I was also supported by an NSF VIGRE fellowship before the start of this project.

DEDICATION

To all those who supported me through the long and trying process of obtaining a doctoral degree, particularly my parents, Nancy and Craig Lemoine.

And to the baristas of the independent coffee shops of Seattle, most notably Fremont Coffee Company. You are the underappreciated heroes of many a student, including myself.

Chapter 1

INTRODUCTION

1.1 Motivation

This work was originally motivated as an extension of the work of Doctor Kirsten Fagnan [33, 34, 35] on finite volume modeling of extracorporeal shock wave therapy. This therapy sends a series of focused, high-amplitude (up to 100 MPa peak pressure) shock waves into the body; it was originally developed to break apart kidney stones without surgery, but has since come into use as an FDA-approved treatment for soft tissue maladies such as plantar fasciitis and tennis elbow, and an experimental treatment for bone conditions such as nonunions — broken bones that do not heal of their own accord — or avascular necrosis of the femoral head, where the head of the femur bone loses blood supply, dies, and collapses.

Doctor Fagnan’s work modeled bone as a nonporous, isotropic elastic solid. This work commenced as an effort to extend her work with a more accurate bone model — bone is often not isotropic, and also contains a hierarchy of multiple different scales of porosity [25], the movement of fluid through which may influence bone growth and remodeling [45]. After a fortuitous meeting at the Institute for Mathematics and its Applications in Minnesota with Professor M. Yvonne Ou, currently of the University of Delaware, the author turned to poroelasticity theory as an improved model of bone.

Poroelasticity theory is a homogenized model for solid porous media containing fluids that can flow through the pore structure. This field was pioneered by Maurice A. Biot, who developed his theory of poroelasticity from the 1930s through the 1960s; a summary of much of Biot’s work can be found in his 1956 and 1962 papers [7, 8, 9]. Biot theory uses linear elasticity to describe the solid portion of the medium (often termed the *skeleton* or *matrix*), linearized compressible fluid dynamics to describe the fluid portion, and Darcy’s law to model the aggregate motion of the fluid through the matrix. While it was originally developed to model fluid-saturated rock and soil, Biot theory has also been used in underwater acoustics [14, 40, 41], and to describe wave propagation in *in vivo* bone [25, 26, 39].

Biot theory predicts rich and complex wave phenomena within poroelastic materials. Three different types of waves appear: fast P waves analogous to standard elastic P waves, in which the fluid and matrix show little relative motion, and typically compress or expand in phase with each other; shear waves analogous to elastic S waves; and slow P waves, where the fluid expands while the solid contracts, or vice versa. The slow P waves exhibit substantial relative motion between the solid and fluid compared to waves of the other two types. The viscosity of the fluid dissipates poroelastic waves as they propagate through the medium, with the fast P and S waves being lightly damped and the slow P wave strongly damped. The viscous dissipation also causes slight dispersion in the fast P and S waves, and strong dispersion in the slow P wave.

Over time, the course of this research has shifted, from being specifically about modeling bone to being more generally about development of finite volume numerical methods for poroelasticity. While in the end there is fairly little material on bone modeling in this thesis, there is a substantial quantity of material on computational methods for poroelasticity and poroelastic-fluid systems in the time domain, including advances in high-resolution finite volume modeling techniques that can be applied more broadly to other problems. The modeling techniques and simulation code developed here also remain available for a return to more bone-specific research at a later date.

1.2 Previous work in computational poroelasticity

A variety of different numerical approaches have been used to model poroelasticity. Carcione, Morency, and Santos provide a thorough review of the previous literature [19]. The earliest numerical work in poroelasticity seems to be that of Garg [37], using a finite difference method in one dimension. Finite difference and pseudospectral methods have continued to be popular since then, with further work by Mikhailenko [62], Hassanzadeh [44], Dai et al. [28], and more recently Chiavassa and Lombard [23], among others. Finite element approaches began being used in the 1980s, with Santos and Oreña’s work [68] being one of the first. Boundary element methods have also been used, such as in the work of Attenborough, Berry, and Chen [4]. Spectral element methods have been used in both the frequency domain [30] and the time domain [63]. Naumovich [65] used a staggered grid finite volume approach in three dimensions for isotropic poroelasticity on a rectilinear grid, with a semidiscrete approach to time-integration. With the recent rise of discontinuous Galerkin methods, DG has been applied to poroelasticity in several works, such as that of de la Puente et al. [29]. There have also been semi-analytical approaches to solving the poroelasticity equations, such as that of Detournay and Cheng [32], who analytically obtain a solution in the Laplace transform domain, but are forced to use an approximate inversion procedure to return to the time domain. Finally, there has been significant work on inverse problems in poroelasticity, for which various forward solvers have been used; of particular note is the paper of Buchanan, Gilbert, and Khashanah [13], who used the finite element method (specifically the FEMLAB software package) to obtain time-harmonic solutions for cancellous bone as part of an inversion scheme to estimate poroelastic material parameters, and the later papers of Buchanan and Gilbert [11, 12], where the authors instead used numerical contour integration of the Green’s function. Numerical work in the 1970s and 1980s focused on isotropic poroelasticity, with the earliest work on anisotropic poroelasticity being by Carcione in 1996 [17].

A major theme in time-domain numerical modeling of poroelasticity has been the difficulty of handling the viscous dissipation term, which has its own intrinsic time scale and causes the poroelasticity system to be stiff, at least if low-frequency waves are being considered. (The time scales associated with dissipation are independent of frequency, so at higher frequencies there is less separation between them and the time scale associated with wave motion — in other words, the system is not stiff if sufficiently high-frequency waves are being considered.) This viscous dissipation is particularly problematic for the slow P wave. While it can still be addressed if Biot’s equations are solved as a unified system (e.g. with

an implicit time-integration method), since the viscous dissipation term is easy to solve analytically in isolation, operator splitting approaches have also been a popular way to address this issue; this is the approach adopted here. Carcione and Quiroga-Goode used operator splitting in conjunction with a pseudospectral method [20], and Chiavassa and Lombard use it with a finite difference approach [23]. De la Puente et al. also investigate an operator splitting approach, and encounter difficulties in obtaining a fast rate of convergence due to the stiffness of the dissipation term, which pushed the problems they investigated toward the diffusive limit of the Biot system [29].

1.3 Contribution of this work

This work solves a velocity-stress formulation of Biot linear orthotropic poroelasticity theory in two and three dimensions using high-resolution finite volume methods on logically rectangular grids. These methods are memory-efficient explicit techniques designed to model hyperbolic systems, and can include wave limiters designed to reduce the effect of numerical dispersion that can cause artifacts in the solution. Since they are based on solving Riemann problems at grid interfaces, it is also straightforward to include material inhomogeneities and interface conditions between cells using these methods. This aids incorporation of interfaces between poroelastic media and fluids into the models considered here. In addition, while this work deals only with logically rectangular grids, finite volume methods lend themselves well to mapping approaches, where the grid is distorted to match the grid lines to the problem geometry. This mapped grid approach is used here to model problems with fluid-poroelastic interfaces that are not straight.

To the author’s knowledge, this is the first use of finite volume methods to model anisotropic poroelasticity, and the first use of high-resolution or mapped-grid finite volume methods to model any kind of poroelasticity. In addition, several novel ideas are developed here to handle, or at least characterize, some of the numerical difficulties posed by poroelasticity:

- Introduction of the energy norm in Section 2.1.6, which allows comparison of stress and velocity variables in a physically meaningful manner, which is useful for measuring error and overall solution magnitude.
- Formulation of the unified treatment of Section 2.4 of closed, open, and imperfect pore conditions at interfaces between poroelastic media and either fluids or other poroelastic media, which allows a single numerical formulation to handle all three cases without difficulty.
- Identification in sections 3.3 (for two dimensions) and 8.4 (for three dimensions) of a subcharacteristic condition that holds between the wave speeds of the full poroelastic system and the reduced system toward which it tends due to the stiff dissipation term. The fact that this subcharacteristic condition is satisfied suggests that solutions obtained via operator splitting will be stable and display the correct wave speeds, despite any other numerical difficulties that may be experienced.

- Identification of a stiff regime in Section 4.3, where the solution time step is long enough that convergence rates are reduced to first-order even when using Strang splitting, which is nominally second-order, for the dissipative term.
- Development of a solution technique for the wave speeds and eigenvectors in sections 5.2.1 and 8.3.1 that allows efficient enough computation that the eigensolver be embedded in the Riemann solver, providing high-accuracy solutions for every Riemann problem.
- Construction in Section 5.2.3 of a method to solve Riemann problems involving nearly arbitrary interface conditions, including Riemann problems posed between domains governed by different PDEs.
- Construction in Section 5.3.1 of a method for transverse solution propagation between dissimilar media that respects the underlying PDE governing each medium.
- Formulation of a revised limiter algorithm in Section 8.5 constructed to correctly limit shear waves in three-dimensional poroelasticity.

1.4 *Outline*

The contents of this thesis are organized as follows: Chapter 2 covers the governing partial differential equations of orthotropic poroelasticity in two and three dimensions. Because poroelastic materials often occur in company with fluids, this chapter also briefly discusses the linear acoustic theory used to model waves in fluids. Interfaces between different media — either between two poroelastic media or between a poroelastic medium and a fluid — are often found in practical problems, so this chapter also formulates interface conditions that relate the solutions of the equations governing adjacent parts of the problem domain.

Following this comes the numerical work of the thesis, which divides roughly into three successive bodies of work, each building on the previous. The first consists of chapters 3 and 4, covering the initial development of a numerical solution procedure for 2D poroelasticity on rectilinear grids in the CLAWPACK finite volume method framework. Chapter 3 lays out the formulation of the numerical method. Because the governing PDEs for poroelasticity form a hyperbolic system with a stiff relaxation term, which is known to cause difficulties for the type of numerical method used here [24, 58], this chapter also includes some heuristic considerations for whether the numerical solution should be expected to be correct. Following this, Chapter 4 covers numerical verification results for poroelasticity on rectilinear grids, against both known analytical solutions and other numerical solutions from the poroelasticity literature. These two chapters correspond to the numerical methods and results sections of the author’s paper with Ou and LeVeque [56].

The second body of work consists of chapters 5-7, which discuss extension of these rectilinear grid solution techniques for all-poroelastic systems to logically-rectangular mapped grid models of systems containing both poroelastic media and fluids. Chapter 5 details

the extension of the numerical method of Chapter 3, including solution of Riemann problems involving orthotropic poroelastic media at grid faces oriented in arbitrary directions, and practical implementation of interface conditions in the Riemann solution process. This chapter also covers modifications necessary to classical high-resolution finite volume methods in order to obtain qualitatively correct solutions in fluid-poroelastic systems, or other such systems where different governing equations may hold on either side of a grid interface. Chapter 6 then verifies the resulting numerical code on a rectilinear grid against analytical solutions for reflection and transmission of plane waves at a planar interface between a fluid and a poroelastic medium, or between two poroelastic media; Chapter 7 extends this to mapped grids with verification against an analytical solution for a plane acoustic wave train scattered by an isotropic poroelastic cylinder, and also models an acoustic pulse striking a simplified model of a human femur. The work of these three chapters is closely analogous to that of Chiavassa and Lombard [22, 23], but using a finite volume method as opposed to their finite difference method.

The third and final body of work consists of chapters 8 and 9, which extend the preceding work to three dimensions. In Chapter 8 the mapped-grid numerical method is extended to three dimensions, and issues that are unique to 3D problems are addressed, such as proper limiting of shear waves in anisotropic materials on a mapped grid. Chapter 9 then applies these methods to some simple three-dimensional problems. Finally, Chapter 10 summarizes the work that has been done here, and lists some of the many opportunities for future work building upon this thesis.

Chapter 2

GOVERNING EQUATIONS

This chapter covers the governing partial differential equations used to model the systems of poroelastic and fluid media treated here, along with related machinery such as a strictly convex entropy function for the poroelastic part of the system, and interface conditions that relate the solutions of the respective PDEs at boundaries between different media. Section 2.1 discusses two-dimensional poroelasticity theory, Section 2.2 covers poroelasticity in three dimensions, Section 2.3 briefly gives the formulations of linear acoustics used here in two and three dimensions, and Section 2.4 discusses interface conditions. This chapter only concerns the formulation of PDE models; discussion of numerical methods for solving them is deferred to chapters 3, 5, and 8.

2.1 *Two dimensional orthotropic, transversely isotropic poroelasticity*

Biot's equations of poroelasticity are a complicated system of PDEs that exhibit rich and varied behaviors. The reader is encouraged to refer to a detailed treatment of the subject, such as Chapter 7 of Carcione [18], but we also provide an overview here. After a review of the basic relations modeling the behavior of a poroelastic medium in sections 2.1.1 through 2.1.4, Section 2.1.5 presents the linear first-order system of PDEs that forms the basis of the numerical model. Section 2.1.6 then defines an energy norm for the state vector that solves some of the scaling issues associated with modeling poroelasticity in SI units. The matrix created for this energy norm allows exploration of some useful properties of the model: Section 2.1.7 provides a concise proof that the first-order system is hyperbolic, and Section 2.1.8 exhibits a strictly convex entropy function for the system, which has implications for the correctness of the numerical solutions, which are explored later in Section 3.3. The reader should note that the formulation of the governing PDEs of poroelasticity in two dimensions as a first-order hyperbolic system (sections 2.1.1 through 2.1.5) is primarily the work of Professor M. Yvonne Ou of the University of Delaware, and that this section corresponds to Section 2 of the author's paper with Ou and LeVeque [56].

2.1.1 *Stress-strain relation*

This work assumes that the constituent material of the solid matrix is isotropic and the anisotropy of the solid matrix results purely from the pore-level microstructure. It also assumes that the anisotropy has a specific form — that the medium is orthotropic, possessing three orthogonal planes of symmetry, and is isotropic with respect to some axis of symmetry. This type of anisotropy is common in engineering composites [38], and in biological materials [27], as well as being present in certain types of stone. Let the z -axis be this axis of symmetry. The elastic stiffness tensor \mathbf{C} of such an orthotropic, transversely isotropic

medium contains five independent components. In its principal axes, and using shorthand notation, \mathbf{C} can be arranged as

$$\mathbf{C} = \begin{pmatrix} c_{11} & c_{12} & c_{13} & 0 & 0 & 0 \\ c_{12} & c_{11} & c_{13} & 0 & 0 & 0 \\ c_{13} & c_{13} & c_{33} & 0 & 0 & 0 \\ 0 & 0 & 0 & c_{55} & 0 & 0 \\ 0 & 0 & 0 & 0 & c_{55} & 0 \\ 0 & 0 & 0 & 0 & 0 & \frac{c_{11}-c_{12}}{2} \end{pmatrix}, \quad (2.1)$$

with the stress tensor and engineering strains arranged into vectors $\boldsymbol{\tau}$ and \mathbf{e} of the form

$$\boldsymbol{\tau} = (\tau_{11} \ \tau_{22} \ \tau_{33} \ \tau_{23} \ \tau_{13} \ \tau_{12})^T \quad \text{and} \quad \mathbf{e} = (\epsilon_{11} \ \epsilon_{22} \ \epsilon_{33} \ 2\epsilon_{23} \ 2\epsilon_{13} \ 2\epsilon_{12})^T. \quad (2.2)$$

Note the factor of 2 applied to the shear strains to convert from the tensor strain $\epsilon_{ij} = \frac{1}{2}(\partial_i u_j + \partial_j u_i)$ to the engineering strain. With this shorthand notation, the stress-strain relation is

$$\boldsymbol{\tau} = \mathbf{C}\mathbf{e}. \quad (2.3)$$

2.1.2 Energy densities and the dissipation potential

One useful property of the poroelasticity system is that it admits an energy density, which forms the basis for an energy norm that will be used extensively.

This section is mainly based on Biot's 1956 papers [7, 8] and Chapter 7 of Carcione's book [18]. All formulations are in terms of the following variables:

1. \mathbf{u} , the displacement vector of the solid matrix
2. $\mathbf{w} := \phi(\mathbf{U} - \mathbf{u})$, the relative motion of the fluid scaled by the porosity, where \mathbf{U} is the displacement vector of the pore fluid and ϕ is the porosity of the medium
3. $\zeta := -\nabla \cdot \mathbf{w}$, the variation in fluid content

Strain energy

In terms of the undrained elasticity tensor \mathbf{C}^u and the strain components of the solid matrix $e_{ii}^{(m)} := \partial_i u_i$ and $e_{ij}^{(m)} := \partial_i u_j + \partial_j u_i$, $i \neq j$, the strain energy density of the Biot model for transversely isotropic materials in three dimensions, with z -axis being the axis of symmetry, is given by

$$\begin{aligned} 2V = & c_{11}^u \left(e_{11}^{(m)2} + e_{22}^{(m)2} \right) + c_{33}^u e_{33}^{(m)2} + 2c_{12}^u e_{11}^{(m)} e_{22}^{(m)} + 2c_{13}^u \left(e_{11}^{(m)} + e_{22}^{(m)} \right) e_{33}^{(m)} \\ & + c_{55}^u (e_{23}^{(m)2} + e_{13}^{(m)2}) + c_{66}^u e_{12}^{(m)2} - 2\alpha_1 M \left(e_{11}^{(m)} + e_{22}^{(m)} \right) \zeta - 2\alpha_3 M e_{33}^{(m)} \zeta + M \zeta^2, \quad (2.4) \end{aligned}$$

where $c_{66}^u = \frac{c_{11}^u - c_{12}^u}{2}$ for a transversely isotropic material and the undrained elastic constants c_{ij}^u are related to those of the drained matrix, c_{ij} , via

$$c_{ij}^u = c_{ij} + M\alpha_i\alpha_j, \quad i, j = 1, \dots, 6 \quad (2.5)$$

$$\boldsymbol{\alpha} := (\alpha_1, \alpha_1, \alpha_3, 0, 0, 0) \quad (2.6)$$

$$\alpha_1 := 1 - \frac{c_{11} + c_{12} + c_{13}}{3K_s} \quad (2.7)$$

$$\alpha_3 := 1 - \frac{2c_{13} + c_{33}}{3K_s} \quad (2.8)$$

$$M := \frac{K_s^2}{K_s [1 + \phi(K_s/K_f - 1)] - (2c_{11} + c_{33} + 2c_{12} + 4c_{13})/9}. \quad (2.9)$$

Here K_s and K_f are the bulk moduli of the constituent material of the solid matrix and of the pore fluid, respectively. The total stresses (solid matrix plus pore pressure) acting on a volume element of the medium are given by the derivatives of strain energy with respect to the associated strains,

$$\tau_{ij} = \frac{\partial V}{\partial e_{ij}^{(m)}}, \quad (2.10)$$

or, in short notation,

$$\tau_I = \sum_{J=1}^6 c_{IJ}^u e_J^{(m)} - M\alpha_I \zeta. \quad (2.11)$$

Similarly, the pore pressure p is

$$p = \frac{\partial V}{\partial \zeta} = M \left(\zeta - \sum_{j=1}^3 \alpha_j e_{jj}^{(m)} \right). \quad (2.12)$$

The two-dimensional work here is for plane-strain conditions in the x - z plane. For these conditions, a function \tilde{V} will later be needed that maps from the pressure and the in-plane stress components τ_{11} , τ_{13} , τ_{33} at a point to the strain energy density at that point. To obtain this, let us first derive the strains as a function of stress for plane-strain conditions by setting $e_{12}^{(m)} = e_{22}^{(m)} = e_{23}^{(m)} = 0$ in equations (2.11) and (2.12), then solving for the remaining strains and the variation in fluid content:

$$\mathbf{e}_{\text{plane}} = \mathbf{S}_{\text{plane}} \boldsymbol{\tau}_{\text{plane}}, \quad (2.13)$$

where the matrices $\mathbf{e}_{\text{plane}}$ and $\boldsymbol{\tau}_{\text{plane}}$ are

$$\mathbf{e}_{\text{plane}} = \left(e_{11}^{(m)} \quad e_{33}^{(m)} \quad e_{13}^{(m)} \quad \zeta \right)^T, \quad \boldsymbol{\tau}_{\text{plane}} = \left(\tau_{11} \quad \tau_{33} \quad \tau_{13} \quad p \right)^T, \quad (2.14)$$

and the undrained in-plane compliance matrix $\mathbf{S}_{\text{plane}}$ is

$$\mathbf{S}_{\text{plane}} = \begin{pmatrix} \frac{c_{33}}{c_{11}c_{33} - (c_{13})^2} & -\frac{c_{13}}{c_{11}c_{33} - (c_{13})^2} & 0 & \frac{\alpha_1 c_{33} - \alpha_3 c_{13}}{c_{11}c_{33} - (c_{13})^2} \\ -\frac{c_{13}}{c_{11}c_{33} - (c_{13})^2} & \frac{c_{11}}{c_{11}c_{33} - (c_{13})^2} & 0 & \frac{\alpha_3 c_{11} - \alpha_1 c_{13}}{c_{11}c_{33} - (c_{13})^2} \\ 0 & 0 & \frac{1}{c_{55}} & 0 \\ \frac{\alpha_1 c_{33} - \alpha_3 c_{13}}{c_{11}c_{33} - (c_{13})^2} & \frac{\alpha_3 c_{11} - \alpha_1 c_{13}}{c_{11}c_{33} - (c_{13})^2} & 0 & \frac{1}{M} + \frac{\alpha_1^2 c_{33} + \alpha_3^2 c_{11} - 2\alpha_1 \alpha_3 c_{13}}{c_{11}c_{33} - (c_{13})^2} \end{pmatrix}. \quad (2.15)$$

Substituting (2.13) through (2.15) into (2.4) yields the strain energy as a function of stress for plane strain conditions, which can be expressed as the quadratic form

$$\tilde{V} := \frac{1}{2} \boldsymbol{\tau}_{\text{plane}}^T \mathbf{S}_{\text{plane}} \boldsymbol{\tau}_{\text{plane}} \quad (2.16)$$

Note the use of the drained elastic coefficients (c_{11} , etc.) rather than the undrained coefficients (c_{11}^u). This quadratic form should be positive-definite on physical grounds — if it were not, then it would be possible to deform the medium or change its fluid content without doing work.

Kinetic energy and the dissipation potential

For anisotropic poroelastic media, the kinetic energy density has the form

$$T = \frac{1}{2} \left(\dot{\mathbf{u}}^T \mathbf{P} \dot{\mathbf{u}} + 2\dot{\mathbf{u}}^T \mathbf{R} \dot{\mathbf{U}} + \dot{\mathbf{U}}^T \mathbf{T} \dot{\mathbf{U}} \right), \quad (2.17)$$

where \mathbf{R} is the induced mass matrix. Assume all the three matrices are diagonal in the principal material axes so that $\mathbf{P} = \text{diag}(a_1, a_2, a_3)$, $\mathbf{R} = \text{diag}(r_1, r_2, r_3)$ and $\mathbf{T} = \text{diag}(t_1, t_2, t_3)$; this is reasonable if the mechanical anisotropy comes from anisotropy of the pore structure. By looking at the special case of no relative motion between fluid and solid, it can be shown that the mass coefficients satisfy the two equations (given as (7.169) in Carcione [18])

$$a_i + r_i = (1 - \phi)\rho_s, \quad r_i + t_i = \phi\rho_f. \quad (2.18)$$

Defining the velocity variables

$$\mathbf{v} := \dot{\mathbf{u}}, \quad \mathbf{q} := \dot{\mathbf{w}}, \quad (2.19)$$

the kinetic energy density can be expressed as

$$\begin{aligned} T &= \frac{1}{2} \sum_{i=1}^3 \left[(1 - \phi)\rho_s v_i^2 - r_i (v_i - \dot{U}_i)^2 + \phi\rho_f \dot{U}_i^2 \right] \\ &= \frac{1}{2} \sum_{i=1}^3 \left[\rho v_i^2 + 2\rho_f q_i v_i + \left(\frac{\rho_f \phi - r_i}{\phi^2} \right) q_i^2 \right] =: \tilde{T}(\mathbf{v}, \mathbf{q}), \end{aligned} \quad (2.20)$$

where ρ_s and ρ_f are the constituent solid density and pore fluid density, respectively and $\rho = (1 - \phi)\rho_s + \phi\rho_f$ is the bulk density of the medium. The induced mass parameters r_i are related to the tortuosity T_i by

$$r_i = \phi\rho_f(1 - T_i). \quad (2.21)$$

Assuming the pore fluid flow is of the Poiseuille type, the dissipation potential Φ_D in an anisotropic medium in terms of the dynamic viscosity η of pore fluid and the permeability tensor \mathbf{K} is

$$\Phi_D = \frac{1}{2} \sum_{i,j=1}^3 \eta (\mathbf{K}^{-1})_{ij} q_i q_j. \quad (2.22)$$

Assume that \mathbf{K} has the same principal directions as \mathbf{P} , \mathbf{R} and \mathbf{T} with eigenvalues $\kappa_1, \kappa_2, \kappa_3$. Then we have

$$\Phi_D = \frac{1}{2} \sum_{i=1}^3 \frac{\phi^2 \eta}{\kappa_i} (\dot{U}_i - v_i)^2 = \frac{1}{2} \sum_{i=1}^3 \frac{\eta}{\kappa_i} q_i^2 =: \tilde{\Phi}_D(\mathbf{q}) \quad (2.23)$$

This form of the dissipation potential is valid at frequencies below cutoff frequency ω_c , defined as

$$\omega_c := \min_i \frac{\eta \phi}{\rho_f T_i \kappa_i}. \quad (2.24)$$

2.1.3 Equations of motion

The equations of motion for the solid part are given in terms of the energy densities and dissipation potential as

$$\partial_t \left(\frac{\partial T}{\partial v_i} \right) + \frac{\partial \Phi_D}{\partial v_i} = \sum_{j=1}^3 \partial_j \left(\frac{\partial V}{\partial e_{ij}^{(m)}} + \phi p \delta_{ij} \right), \quad (2.25)$$

or equivalently

$$\partial_t \left(\frac{\partial \tilde{T}}{\partial v_i} - \phi \frac{\partial \tilde{T}}{\partial q_i} \right) - \phi \frac{\partial \tilde{\Phi}_D}{\partial q_i} = \sum_{j=1}^3 \partial_j \left(\frac{\partial V}{\partial e_{ij}^{(m)}} + \phi p \delta_{ij} \right). \quad (2.26)$$

Similarly, the equation of motion for the fluid part is

$$\partial_t \left(\frac{\partial T}{\partial \dot{U}_i} \right) + \frac{\partial \Phi_D}{\partial \dot{U}_i} = -\phi \sum_{j=1}^3 \partial_j \left(\frac{\partial V}{\partial \zeta} \delta_{ij} \right), \quad (2.27)$$

or equivalently

$$\phi \partial_t \left(\frac{\partial \tilde{T}}{\partial q_i} \right) + \phi \frac{\partial \tilde{\Phi}_D}{\partial q_i} = -\phi \sum_{j=1}^3 \partial_j \left(\frac{\partial V}{\partial \zeta} \delta_{ij} \right), \quad (2.28)$$

which reduces to

$$-\partial_i p = \rho_f \dot{v}_i + \left(\frac{\rho_f \phi - r_i}{\phi^2} \right) \dot{q}_i + \frac{\eta}{\kappa_i} q_i, \quad i = 1, 2, 3. \quad (2.29)$$

The equations of motion for the fluid-solid composite are obtained by adding (2.26) with (2.28), giving

$$\sum_{j=1}^3 \partial_j \tau_{ij} = \partial_t \left(\frac{\partial \tilde{T}}{\partial v_i} \right) = \rho_f \dot{v}_i + \rho_f \dot{q}_i. \quad (2.30)$$

2.1.4 Governing equations for plane-strain case

Since the material is assumed isotropic in the x - y plane, let us consider the plane strain problem in the x - z plane. The governing equations are obtained by suppressing the y -component (subscript 2) of \mathbf{U} , \mathbf{u} , \mathbf{w} , \mathbf{q} and those terms of $e_{ij}^{(m)}$ with $i = 2$ or $j = 2$ in V , T , \tilde{T} , Φ_D , and $\tilde{\Phi}_D$. (While nonzero out-of-plane stresses do arise in a plane-strain problem, they do not produce in-plane motion, and can be ignored for purposes of studying the in-plane dynamics of the medium. By definition, for plane strain conditions the out-of-plane strains are zero.) This results in two types of governing equation:

- Stress-strain relations, obtained by differentiating (2.12) and (2.11) for $I = 1, 3, 5$ with respect to time,

$$\partial_t \tau_{xx} = c_{11}^u \partial_x v_x + c_{13}^u \partial_z v_z + \alpha_1 M (\partial_x q_x + \partial_z q_z) + \partial_t s_1 \quad (2.31)$$

$$\partial_t \tau_{zz} = c_{13}^u \partial_x v_x + c_{33}^u \partial_z v_z + \alpha_3 M (\partial_x q_x + \partial_z q_z) + \partial_t s_3 \quad (2.32)$$

$$\partial_t \tau_{xz} = c_{55}^u (\partial_z v_x + \partial_x v_z) + \partial_t s_5 \quad (2.33)$$

$$\partial_t p = -\alpha_1 M \partial_x v_x - \alpha_3 M \partial_z v_z - M (\partial_x q_x + \partial_z q_z) + \partial_t s_f, \quad (2.34)$$

where s_1 , s_3 , s_5 , and s_f are the solid and fluid external sources.

- Equations of motion

$$\rho \partial_t v_x + \rho_f \partial_t q_x = \partial_x \tau_{xx} + \partial_z \tau_{xz} \quad (2.35)$$

$$\rho \partial_t v_z + \rho_f \partial_t q_z = \partial_x \tau_{xz} + \partial_z \tau_{zz} \quad (2.36)$$

$$\rho_f \partial_t v_x + m_1 \partial_t q_x + \left(\frac{\eta}{\kappa_1} \right) q_x = -\partial_x p \quad (2.37)$$

$$\rho_f \partial_t v_z + m_3 \partial_t q_z + \left(\frac{\eta}{\kappa_3} \right) q_z = -\partial_z p, \quad (2.38)$$

where $m_i := \frac{\rho_f \phi - r_i}{\phi^2} = \frac{\rho_f T_i}{\phi}$.

2.1.5 Governing equations as a linear first-order system

Solving (2.35) and (2.37) for $\partial_t v_x$ and $\partial_t q_x$, results in

$$\partial_t v_x = \frac{1}{\Delta_1} \left(m_1 \partial_x \tau_{xx} + m_1 \partial_z \tau_{xz} + \rho_f \partial_x p + \rho_f \frac{\eta}{\kappa_1} q_x \right) \quad (2.39)$$

$$\partial_t q_x = \frac{1}{\Delta_1} \left(-\rho_f \partial_x \tau_{xx} - \rho_f \partial_z \tau_{xz} - \rho \partial_x p - \rho \frac{\eta}{\kappa_1} q_x \right), \quad (2.40)$$

where $\Delta_1 := \rho m_1 - \rho_f^2$. Similarly, (2.36) and (2.38) lead to

$$\partial_t v_z = \frac{1}{\Delta_3} \left(m_3 \partial_x \tau_{xz} + m_3 \partial_z \tau_{zz} + \rho_f \partial_z p + \rho_f \frac{\eta}{\kappa_3} q_z \right) \quad (2.41)$$

$$\partial_t q_z = \frac{1}{\Delta_3} \left(-\rho_f \partial_x \tau_{xz} - \rho_f \partial_z \tau_{zz} - \rho \partial_z p - \rho \frac{\eta}{\kappa_3} q_z \right), \quad (2.42)$$

where $\Delta_3 := \rho m_3 - \rho_f^2$. Combining the stress-strain relations (2.31)-(2.34) with equations (2.39)-(2.42), we obtain the 8×8 linear first-order system

$$\partial_t \mathbf{Q} + \mathbf{A} \partial_x \mathbf{Q} + \mathbf{B} \partial_z \mathbf{Q} = \mathbf{D} \mathbf{Q} + \partial_t \mathbf{s}, \quad (2.43)$$

where

$$\mathbf{Q} = (p \quad \tau_{xx} \quad \tau_{zz} \quad \tau_{xz} \quad v_x \quad v_z \quad q_x \quad q_z)^T, \quad (2.44)$$

$$\mathbf{A} = \begin{pmatrix} 0_{4 \times 4} & \mathbf{A}_{sv} \\ \mathbf{A}_{vs} & 0_{4 \times 4} \end{pmatrix}, \quad \mathbf{B} = \begin{pmatrix} 0_{4 \times 4} & \mathbf{B}_{sv} \\ \mathbf{B}_{vs} & 0_{4 \times 4} \end{pmatrix}, \quad \mathbf{D} = \begin{pmatrix} 0_{4 \times 4} & 0_{4 \times 4} \\ 0_{4 \times 4} & \mathbf{D}_v \end{pmatrix}, \quad (2.45)$$

and

$$\mathbf{s} = (s_f \quad s_1 \quad s_3 \quad s_5 \quad 0 \quad 0 \quad 0 \quad 0)^T. \quad (2.46)$$

The blocks forming matrices \mathbf{A} , \mathbf{B} , and \mathbf{D} are

$$\mathbf{A}_{sv} = \begin{pmatrix} \alpha_1 M & 0 & M & 0 \\ -c_{11}^u & 0 & -\alpha_1 M & 0 \\ -c_{13}^u & 0 & -\alpha_3 M & 0 \\ 0 & -c_{55}^u & 0 & 0 \end{pmatrix} \quad \mathbf{A}_{vs} = \begin{pmatrix} -\frac{\rho_f}{\Delta_1} & -\frac{m_1}{\Delta_1} & 0 & 0 \\ 0 & 0 & 0 & -\frac{m_3}{\Delta_3} \\ \frac{\rho}{\Delta_1} & \frac{\rho_f}{\Delta_1} & 0 & 0 \\ 0 & 0 & 0 & \frac{\rho_f}{\Delta_3} \end{pmatrix} \quad (2.47)$$

$$\mathbf{B}_{sv} = \begin{pmatrix} 0 & \alpha_3 M & 0 & M \\ 0 & -c_{13}^u & 0 & -\alpha_1 M \\ 0 & -c_{33}^u & 0 & -\alpha_3 M \\ -c_{55}^u & 0 & 0 & 0 \end{pmatrix} \quad \mathbf{B}_{vs} = \begin{pmatrix} 0 & 0 & 0 & -\frac{m_1}{\Delta_1} \\ -\frac{\rho_f}{\Delta_3} & 0 & -\frac{m_3}{\Delta_3} & 0 \\ 0 & 0 & 0 & \frac{\rho_f}{\Delta_1} \\ \frac{\rho}{\Delta_3} & 0 & \frac{\rho_f}{\Delta_3} & 0 \end{pmatrix} \quad (2.48)$$

$$\mathbf{D}_v = \begin{pmatrix} 0 & 0 & \frac{\rho_f \eta}{\Delta_1 \kappa_1} & 0 \\ 0 & 0 & 0 & \frac{\rho_f \eta}{\Delta_3 \kappa_3} \\ 0 & 0 & -\frac{\rho \eta}{\Delta_1 \kappa_1} & 0 \\ 0 & 0 & 0 & -\frac{\rho \eta}{\Delta_3 \kappa_3} \end{pmatrix}. \quad (2.49)$$

Here the subscripts s and v denote the stress variables (p and $\boldsymbol{\tau}$) and velocity variables (\mathbf{v} and \mathbf{q}), respectively. Ordering the variables to highlight this block structure emphasizes the underlying physics — gradients of stress produce changes in velocity, and gradients of velocity produce changes in stress — and it also proves useful mathematically.

It is (2.43), and its three-dimensional analog derived in the next section, that form the basis for all numerical work here. Note that while the coefficient matrices \mathbf{A} , \mathbf{B} , and \mathbf{D} are defined in the material principal axes, this system can be extended to model media where the principal axes are different from the global x - z axes through an appropriate transformation of the state variables in \mathbf{Q} that come from vector and tensor quantities, and application of the chain rule in the partial derivatives with respect to the spatial variables. In such cases, the principal directions are labeled as the 1 and 3 axes to distinguish them from the computational x and z axes.

It is worth noting that (2.43) is not just a generic “black box” equation, but one of a very specific type: a first-order hyperbolic system with a stiff relaxation source term. We will see that the system is hyperbolic in Section 2.1.7, and the source term $\mathbf{D} \mathbf{Q}$ shows itself to be of relaxation type by having only zero and negative eigenvalues — in the absence of the spatial

derivative terms, it would cause the solution to decay exponentially toward the null space $\mathcal{N}(\mathbf{D})$, and even with the other terms present we can expect it to keep the solution close to $\mathcal{N}(\mathbf{D})$. (Whether the relaxation term really is stiff depends on the other time scales of the particular problem being solved, but it is stiff for some of the problems considered here.) On the subject of time scales, because \mathbf{D} is extremely sparse, we can immediately read off the eigenvalues associated with dissipation in the 1 and 3 axes — respectively, $-\frac{\rho\eta}{\Delta_1\kappa_1}$ and $-\frac{\rho\eta}{\Delta_3\kappa_3}$ — and so define the characteristic time for decay in each axis as the negative inverse of these eigenvalues,

$$\tau_{d1} := \frac{\Delta_1\kappa_1}{\rho\eta}, \quad \tau_{d3} := \frac{\Delta_3\kappa_3}{\rho\eta}. \quad (2.50)$$

2.1.6 Energy norm

Let $\mathcal{E} := \tilde{T} + \tilde{V}$ be the total mechanical energy per unit volume in a representative element, where \tilde{T} is the kinetic energy function from (2.20), and \tilde{V} is the strain energy function from (2.16). The subsequent analysis will use its Hessian with respect to the state variables in \mathbf{Q} , which is the symmetric matrix

$$\mathbf{E} = \begin{pmatrix} \mathbf{E}_s & 0_{4 \times 4} \\ 0_{4 \times 4} & \mathbf{E}_v \end{pmatrix}, \quad (2.51)$$

where in the principal material axes the diagonal blocks of \mathbf{E} are

$$\mathbf{E}_s = \begin{pmatrix} \frac{1}{M} + \frac{\alpha_1^2 c_{33} + \alpha_3^2 c_{11} - 2\alpha_1 \alpha_3 c_{13}}{c_{11} c_{33} - (c_{13})^2} & \frac{\alpha_1 c_{33} - \alpha_3 c_{13}}{c_{11} c_{33} - (c_{13})^2} & \frac{\alpha_3 c_{11} - \alpha_1 c_{13}}{c_{11} c_{33} - (c_{13})^2} & 0 \\ \frac{\alpha_1 c_{33} - \alpha_3 c_{13}}{c_{11} c_{33} - (c_{13})^2} & \frac{c_{33}}{c_{11} c_{33} - (c_{13})^2} & -\frac{c_{13}}{c_{11} c_{33} - (c_{13})^2} & 0 \\ \frac{\alpha_3 c_{11} - \alpha_1 c_{13}}{c_{11} c_{33} - (c_{13})^2} & -\frac{c_{13}}{c_{11} c_{33} - (c_{13})^2} & \frac{c_{11}}{c_{11} c_{33} - (c_{13})^2} & 0 \\ 0 & 0 & 0 & \frac{1}{c_{55}} \end{pmatrix} \quad (2.52)$$

$$\mathbf{E}_v = \begin{pmatrix} \rho & 0 & \rho_f & 0 \\ 0 & \rho & 0 & \rho_f \\ \rho_f & 0 & m_1 & 0 \\ 0 & \rho_f & 0 & m_3 \end{pmatrix}.$$

The matrix \mathbf{E} is positive-definite because it is the Hessian of the positive-definite quadratic form \mathcal{E} . In fact, the individual blocks \mathbf{E}_s and \mathbf{E}_v are themselves symmetric positive-definite matrices. Because \mathcal{E} has no linear terms in the state variables, it can be written compactly in terms of its Hessian as

$$\mathcal{E} = \frac{1}{2} \mathbf{Q}^T \mathbf{E} \mathbf{Q}. \quad (2.53)$$

For many poroelastic materials, the components of \mathbf{Q} are very badly scaled relative to each other when expressed in common units — for example, waves in geological materials typically have stress components about seven orders of magnitude larger than their velocity components when expressed in SI base units. This makes using the usual vector norms on \mathbf{Q} problematic, but this issue can be fixed by using \mathbf{E} to define an energy norm,

$$\|\mathbf{Q}\|_E := \sqrt{\mathbf{Q}^H \mathbf{E} \mathbf{Q}}. \quad (2.54)$$

Thus $\mathbf{E}\check{\mathbf{A}}$ is a symmetric matrix, and (2.59) is a real symmetric-definite generalized eigenproblem. As such, it has purely real eigenvalues and a full set of linearly independent eigenvectors. The matrix $\check{\mathbf{A}}$ is therefore diagonalizable and has pure real eigenvalues, which means that (2.57) is a hyperbolic system.

An additional useful result can be derived from (2.59). Because it is a symmetric-definite eigenproblem, its eigenvectors are orthogonal with respect to the matrix \mathbf{E} — if \mathbf{v}_i and \mathbf{v}_j correspond to different eigenvalues,

$$\mathbf{v}_i^T \mathbf{E} \mathbf{v}_j = 0. \quad (2.61)$$

This property will become useful starting in Chapter 5.

2.1.8 Entropy function

With some additional effort we can also see that the energy density is a strictly convex entropy function of the system (2.43), in a sense similar to that of Chen, Levermore, and Liu [21]. Adapting the definition of [21] to the notation used here, a function $\Phi : \mathbb{R}^8 \rightarrow \mathbb{R}$ is a strictly convex entropy function for the system (2.43) if it satisfies the following conditions:

1. $\Phi''(\mathbf{Q})(n_x \mathbf{A} + n_z \mathbf{B})$ is symmetric for all scalars n_x and n_z
2. $(\Phi'(\mathbf{Q}))^T \mathbf{D} \mathbf{Q} \leq 0$ for all $\mathbf{Q} \in \mathbb{R}^8$
3. For $\mathbf{Q} \in \mathbb{R}^8$, the following are equivalent:
 - (a) $\mathbf{D} \mathbf{Q} = 0$
 - (b) $(\Phi'(\mathbf{Q}))^T \mathbf{D} \mathbf{Q} = 0$
4. $\Phi''(\mathbf{Q})$ is positive-definite

Here the primes indicate gradients with respect to \mathbf{Q} , so Φ'' is the Hessian of Φ with respect to \mathbf{Q} . The definition of Chen, Levermore, and Liu includes an additional clause in item 3 related to an operator called $\mathbf{\Pi}$ here (\mathcal{Q} in their notation) that maps from \mathbf{Q} to the conserved quantities of the relaxation part of the system, $\partial_t \mathbf{Q} = \mathbf{D} \mathbf{Q}$ — namely, that conditions 3(a) and 3(b) should also be equivalent to $\Phi'(\mathbf{Q})^T = \mathbf{v}^T \mathbf{\Pi}$ for some appropriately-sized vector \mathbf{v} . Rather than take this as part of the definition of a strictly convex entropy function, it is more convenient here to take it as a requirement on $\mathbf{\Pi}$.

Suppose $\Phi(\mathbf{Q}) = \mathcal{E}(\mathbf{Q}) = \frac{1}{2} \mathbf{Q}^T \mathbf{E} \mathbf{Q}$. From the preceding sections we already know that conditions 1 and 4 are satisfied, so it only remains to prove conditions 2 and 3. Since $\Phi'(\mathbf{Q}) = \mathbf{E} \mathbf{Q}$, condition 2 reduces to $\mathbf{Q}^T \mathbf{E} \mathbf{D} \mathbf{Q} \leq 0$ for all $\mathbf{Q} \in \mathbb{R}^8$. Using (2.45) and (2.51)

for \mathbf{D} and \mathbf{E} , and recalling $\Delta_i = \rho m_i - \rho_f^2$, we find that

$$\mathbf{ED} = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & -\eta/\kappa_1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & -\eta/\kappa_3 & 0 \end{pmatrix}. \quad (2.62)$$

By inspection, \mathbf{ED} is a symmetric negative-semidefinite matrix, so $\Phi'(\mathbf{Q})^T \mathbf{DQ} \leq 0$ for all $\mathbf{Q} \in \mathbb{R}^8$ and condition 2 is satisfied.

For condition 3, note that 3(a) implies 3(b) since if $\mathbf{DQ} = 0$, necessarily $\Phi'(\mathbf{Q})^T \mathbf{DQ} = 0$. To see that $\Phi'(\mathbf{Q})^T \mathbf{DQ} = \mathbf{Q}^T \mathbf{EDQ} = 0$ implies $\mathbf{DQ} = 0$, note that from (2.62), we have $\mathbf{Q}^T \mathbf{EDQ} = -\frac{\eta}{\kappa_1} q_x^2 - \frac{\eta}{\kappa_3} q_z^2 = 0$ if and only if $q_x = q_z = 0$. Since \mathbf{D} only has nonzero entries in the columns corresponding to q_x and q_z , $\Phi'(\mathbf{Q})^T \mathbf{DQ} = 0$ if and only if $\mathbf{DQ} = 0$. Therefore condition 3 holds, and \mathcal{E} is a strictly convex entropy function for the poroelastic system (2.43).

2.1.9 Equipartition of energy

A useful property that can be derived from the block structure of the system is that a simple plane wave of the homogeneous hyperbolic PDE (2.57) carries equal amounts of kinetic and strain energy. To see this, note that a simple traveling plane wave solution has the form $\mathbf{Q}(x, z, t) = \mathbf{r}f(n_x x + n_z z - \lambda t)$, where \mathbf{r} and λ satisfy the eigenproblem

$$\check{\mathbf{A}}\mathbf{r} = \lambda\mathbf{r}, \quad \check{\mathbf{A}} := n_x \mathbf{A} + n_z \mathbf{B}. \quad (2.63)$$

For a traveling wave, we also know that λ is nonzero. Dividing \mathbf{r} into stress and velocity parts according to the block partitioning of the system, $\mathbf{r} =: (\mathbf{r}_s^T \quad \mathbf{r}_v^T)^T$, this becomes a pair of equations,

$$\check{\mathbf{A}}_{sv} \mathbf{r}_v = \lambda \mathbf{r}_s, \quad \check{\mathbf{A}}_{vs} \mathbf{r}_s = \lambda \mathbf{r}_v, \quad (2.64)$$

or, multiplying through by \mathbf{E} ,

$$\mathbf{E}_s \check{\mathbf{A}}_{sv} \mathbf{r}_v = \lambda \mathbf{E}_s \mathbf{r}_s, \quad \mathbf{E}_v \check{\mathbf{A}}_{vs} \mathbf{r}_s = \lambda \mathbf{E}_v \mathbf{r}_v. \quad (2.65)$$

Now, multiply the first equation of (2.65) from the left by \mathbf{r}_s^T , giving

$$\mathbf{r}_s^T \mathbf{E}_s \check{\mathbf{A}}_{sv} \mathbf{r}_v = \lambda \mathbf{r}_s^T \mathbf{E}_s \mathbf{r}_s, \quad (2.66)$$

and notice the symmetry relation between blocks, which holds because \mathbf{EA} and \mathbf{EB} are symmetric:

$$\mathbf{E}_s \mathbf{A}_{sv} = (\mathbf{E}_v \mathbf{A}_{vs})^T, \quad \mathbf{E}_s \mathbf{B}_{sv} = (\mathbf{E}_v \mathbf{B}_{vs})^T. \quad (2.67)$$

Using this symmetry relation and the eigenproblem (2.65), $\mathbf{r}_s^T \mathbf{E}_s \check{\mathbf{A}}_{sv} = \mathbf{r}_s^T (\mathbf{E}_v \check{\mathbf{A}}_{vs})^T = (\mathbf{E}_v \check{\mathbf{A}}_{vs} \mathbf{r}_s)^T = (\lambda \mathbf{E}_v \mathbf{r}_v)^T$. Since λ is nonzero for a traveling wave, it may be divided out; remembering that \mathbf{E}_v is symmetric, we obtain

$$\mathbf{r}_v^T \mathbf{E}_v \mathbf{r}_v = \mathbf{r}_s \mathbf{E}_s \mathbf{r}_s. \quad (2.68)$$

Note that there is no need to distinguish between the ordinary transpose and the Hermitian here, since as demonstrated earlier the eigenvectors are real-valued. Note also that there was no specific reference to the physics of poroelastic medium here, only to general algebraic properties of the system — this property therefore holds for the three-dimensional poroelasticity system developed in the next section, and for linear acoustics, discussed in Section 2.3, because they satisfy the same block symmetry relations.

2.2 Three dimensional orthotropic poroelasticity

2.2.1 Stress rate-velocity relations

Equations (7.131) and (7.133) of Carcione [18] give the stress-strain relation for an anisotropic poroelastic material in an orthogonal set of axes labeled 1, 2, and 3. Using the summation convention for repeated indices, these equations are

$$\begin{aligned} p &= M(\zeta - \alpha_I e_I) \\ \tau_I &= c_{IJ}^u e_J - M \alpha_I \zeta. \end{aligned} \quad (2.69)$$

The quantities in this system are defined as follows:

- p is the fluid pressure
- ζ is the variation of fluid content, $\zeta = -\text{div}(\phi(\mathbf{u}_f - \mathbf{u}_m))$
- ϕ is the porosity of the material
- \mathbf{u}_f and \mathbf{u}_m are the displacements of the fluid and solid, respectively, from their stress-free configurations
- e_I is the I 'th component of engineering strain, ordered as $e = (\epsilon_{11} \ \epsilon_{22} \ \epsilon_{33} \ \gamma_{23} \ \gamma_{13} \ \gamma_{12})^T$. Note that the engineering shear strains γ_{ij} are twice the tensor shear strains ϵ_{ij} , $i \neq j$.
- τ_I is the I 'th component of the total stress in the material, taken in the same order as the strain
- $c_{IJ}^u = c_{IJ} + \alpha_I \alpha_J M$ is the undrained elastic stiffness tensor of the matrix
- c_{IJ} is the drained elastic stiffness tensor of the matrix
- α_I is the I 'th effective stress coefficient, given by $\alpha_I = 1 - \frac{1}{3K_s} \sum_{J=1}^3 c_{IJ}$

- K_s is the bulk modulus of the matrix material
- M is a parameter related to the bulk compressibility of the medium,

$$M = K_s \left(\left(1 - \frac{K^*}{K_s} \right) - \phi \left(1 - \frac{K_s}{K_f} \right) \right)^{-1}$$

- K_f is the bulk modulus of the fluid
- K^* is another bulk stiffness coefficient, $K^* = \frac{1}{9} \sum_{I=1}^3 \sum_{J=1}^3 c_{IJ}$

To cast this in the form needed to build a first-order velocity-stress system, note the following relations between the velocities and strain rates (presuming that space and time derivatives can be interchanged):

$$\begin{aligned} \partial_t e_1 &= v_{1,1} \\ \partial_t e_2 &= v_{2,2} \\ \partial_t e_3 &= v_{3,3} \\ \partial_t e_4 &= v_{2,3} + v_{3,2} \\ \partial_t e_5 &= v_{1,3} + v_{3,1} \\ \partial_t e_6 &= v_{1,2} + v_{2,1} \\ \partial_t \zeta &= -q_{1,1} - q_{2,2} - q_{3,3}. \end{aligned} \tag{2.70}$$

Here \mathbf{v} is the velocity of the matrix relative to an inertial frame, and \mathbf{q} is the flow rate of the fluid relative to the matrix (the porosity ϕ times the aggregate velocity of the fluid relative to the matrix). Subscript indices before a comma represent components, while those after a comma represent differentiation with respect to distance along an axis. Differentiating (2.69) with respect to time, and defining the vectors of stresses and velocities as $\mathbf{Q}_s = (\tau_{11} \tau_{22} \tau_{33} \tau_{23} \tau_{13} \tau_{12} p)^T$ and $\mathbf{Q}_v = (v_1 v_2 v_3 q_1 q_2 q_3)^T$, results in a system relating $\partial_t \mathbf{Q}_s$ to the gradients of the velocities:

$$\partial_t \mathbf{Q}_s + \mathbf{A}_{sv} \mathbf{Q}_{v,1} + \mathbf{B}_{sv} \mathbf{Q}_{v,2} + \mathbf{C}_{sv} \mathbf{Q}_{v,3} = 0. \tag{2.71}$$

Here the pressure is placed at the end of the stress vector so that the first six entries can have the same indices as in the conventional solid mechanics notation, which saves confusion in implementation.

Rather than give the matrices \mathbf{A}_{sv} , \mathbf{B}_{sv} , and \mathbf{C}_{sv} individually, it is more convenient to write the matrix $\check{\mathbf{A}}_{sv} = n_1 \mathbf{A}_{sv} + n_2 \mathbf{B}_{sv} + n_3 \mathbf{C}_{sv}$ that is the coefficient of the directional derivative of \mathbf{Q}_v in the (n_1, n_2, n_3) direction. If the medium is orthotropic, and the 1-2-3 axes are its principal axes, then

$$\check{\mathbf{A}}_{sv} = - \begin{pmatrix} n_1 c_{11}^u & n_2 c_{12}^u & n_3 c_{13}^u & n_1 \alpha_1 M & n_2 \alpha_1 M & n_3 \alpha_1 M \\ n_1 c_{12}^u & n_2 c_{22}^u & n_3 c_{23}^u & n_1 \alpha_2 M & n_2 \alpha_2 M & n_3 \alpha_2 M \\ n_1 c_{13}^u & n_2 c_{23}^u & n_3 c_{33}^u & n_1 \alpha_3 M & n_2 \alpha_3 M & n_3 \alpha_3 M \\ 0 & n_3 c_{44} & n_2 c_{44} & 0 & 0 & 0 \\ n_3 c_{55} & 0 & n_1 c_{55} & 0 & 0 & 0 \\ n_2 c_{66} & n_1 c_{66} & 0 & 0 & 0 & 0 \\ -n_1 M \alpha_1 & -n_2 M \alpha_2 & -n_3 M \alpha_3 & -n_1 M & -n_2 M & -n_3 M \end{pmatrix}. \tag{2.72}$$

2.2.2 Equations of motion

System (2.71) is useful, but does not yet provide a closed set of equations that can be used to describe the dynamics of the poroelastic medium. Equations of motion are still required that relate accelerations to gradients of stress. Equations (7.255) and (7.256) of [18] provide the key. If the medium is orthotropic and the 1-2-3 axes are its principal axes, these equations relate accelerations to stress gradients:

$$\begin{aligned}\tau_{ij,j} &= \rho \partial_t^2 u_{mi} + \rho_f \partial_t^2 w_i \\ -p_{,i} &= \rho_f \partial_t^2 u_{mi} + m_i \partial_t^2 w_i + \frac{\eta}{\kappa_i} \partial_t w_i.\end{aligned}\tag{2.73}$$

Here i ranges from 1 to 3; there is a sum over j in the first equation, but no sum over i in the second. The new variables in these equations are as follows:

- ρ is the bulk density of the medium, $\rho = (1 - \phi)\rho_s + \phi\rho_f$
- ρ_s is the density of the matrix material
- ρ_f is the density of the fluid
- \mathbf{w} is the displacement of the fluid relative to the matrix, scaled by the porosity. The rate of change of \mathbf{w} is $\partial_t \mathbf{w} = \mathbf{q}$.
- m_i is the fluid inertia along axis i , $m_i = \rho_f T_i / \phi$
- T_i is the tortuosity of the matrix along axis i , defined as the factor by which the kinetic energy of the fluid must be higher than its density would indicate for straight-line motion, in order to have a given bulk velocity along that axis

For each i , (2.73) is a system of two equations in two unknowns. Noting that $\partial_t u_{mi} = v_i$, this system becomes

$$\begin{aligned}\rho \partial_t v_i + \rho_f \partial_t q_i &= \tau_{ij,j} \\ \rho_f \partial_t v_i + m_i \partial_t q_i &= -p_{,i} - \frac{\eta}{\kappa_i} q_i.\end{aligned}\tag{2.74}$$

Solving with Cramer's Rule results in

$$\begin{aligned}\partial_t v_i &= \frac{m_i}{\Delta_i} \tau_{ij,j} + \frac{\rho_f}{\Delta_i} p_{,i} + \frac{\rho_f \eta}{\Delta_i \kappa_i} q_i \\ \partial_t q_i &= -\frac{\rho_f}{\Delta_i} \tau_{ij,j} - \frac{\rho}{\Delta_i} p_{,i} - \frac{\rho \eta}{\Delta_i \kappa_i} q_i,\end{aligned}\tag{2.75}$$

where $\Delta_i := \rho m_i - \rho_f^2$.

It is now possible to write a linear system relating the rates of change of the velocities to the gradients of stress, of the form

$$\partial_t \mathbf{Q}_v + \mathbf{A}_{vs} \mathbf{Q}_{s,1} + \mathbf{B}_{vs} \mathbf{Q}_{s,2} + \mathbf{C}_{vs} \mathbf{Q}_{s,3} = \mathbf{D}_v \mathbf{Q}_v.\tag{2.76}$$

Again, it is more convenient to provide $\check{\mathbf{A}}_{vs} = n_1 \mathbf{A}_{vs} + n_2 \mathbf{B}_{vs} + n_3 \mathbf{C}_{vs}$, rather than the individual matrices of (2.76):

$$\check{\mathbf{A}}_{vs} = - \begin{pmatrix} n_1 \frac{m_1}{\Delta_1} & 0 & 0 & 0 & n_3 \frac{m_1}{\Delta_1} & n_2 \frac{m_1}{\Delta_1} & n_1 \frac{\rho_f}{\Delta_1} \\ 0 & n_2 \frac{m_2}{\Delta_2} & 0 & n_3 \frac{m_2}{\Delta_2} & 0 & n_1 \frac{m_2}{\Delta_2} & n_2 \frac{\rho_f}{\Delta_2} \\ 0 & 0 & n_3 \frac{m_3}{\Delta_3} & n_2 \frac{m_3}{\Delta_3} & n_1 \frac{m_3}{\Delta_3} & 0 & n_3 \frac{\rho_f}{\Delta_3} \\ -n_1 \frac{\rho_f}{\Delta_1} & 0 & 0 & 0 & -n_3 \frac{\rho_f}{\Delta_1} & -n_2 \frac{\rho_f}{\Delta_1} & -n_1 \frac{\rho}{\Delta_1} \\ 0 & -n_2 \frac{\rho_f}{\Delta_2} & 0 & -n_3 \frac{\rho_f}{\Delta_2} & 0 & -n_1 \frac{\rho_f}{\Delta_2} & -n_2 \frac{\rho}{\Delta_2} \\ 0 & 0 & -n_3 \frac{\rho_f}{\Delta_3} & -n_2 \frac{\rho_f}{\Delta_3} & -n_1 \frac{\rho_f}{\Delta_3} & 0 & -n_3 \frac{\rho}{\Delta_3} \end{pmatrix}. \quad (2.77)$$

The matrix \mathbf{D}_v models the viscous dissipation, and is given by

$$\mathbf{D}_v = \begin{pmatrix} 0 & 0 & 0 & \frac{\rho_f \eta}{\Delta_1 \kappa_1} & 0 & 0 \\ 0 & 0 & 0 & 0 & \frac{\rho_f \eta}{\Delta_2 \kappa_2} & 0 \\ 0 & 0 & 0 & 0 & 0 & \frac{\rho_f \eta}{\Delta_3 \kappa_3} \\ 0 & 0 & 0 & -\frac{\rho \eta}{\Delta_1 \kappa_1} & 0 & 0 \\ 0 & 0 & 0 & 0 & -\frac{\rho \eta}{\Delta_2 \kappa_2} & 0 \\ 0 & 0 & 0 & 0 & 0 & -\frac{\rho \eta}{\Delta_3 \kappa_3} \end{pmatrix}. \quad (2.78)$$

2.2.3 First-order velocity-stress system

Combining (2.71) and (2.76), and defining the full 13-element state vector as $\mathbf{Q} = (\mathbf{Q}_s^T \mathbf{Q}_v^T)^T$, the first-order stress-velocity system describing three-dimensional poroelasticity is

$$\partial_t \mathbf{Q} + \mathbf{A} \mathbf{Q}_{,1} + \mathbf{B} \mathbf{Q}_{,2} + \mathbf{C} \mathbf{Q}_{,3} = \mathbf{D} \mathbf{Q}, \quad (2.79)$$

where

$$\mathbf{A} = \begin{pmatrix} 0_{7 \times 7} & \mathbf{A}_{sv} \\ \mathbf{A}_{vs} & 0_{6 \times 6} \end{pmatrix}, \quad \mathbf{B} = \begin{pmatrix} 0_{7 \times 7} & \mathbf{B}_{sv} \\ \mathbf{B}_{vs} & 0_{6 \times 6} \end{pmatrix}, \quad \mathbf{C} = \begin{pmatrix} 0_{7 \times 7} & \mathbf{C}_{sv} \\ \mathbf{C}_{vs} & 0_{6 \times 6} \end{pmatrix}, \quad \mathbf{D} = \begin{pmatrix} 0_{7 \times 7} & 0_{6 \times 7} \\ 0_{7 \times 6} & \mathbf{D}_v \end{pmatrix}. \quad (2.80)$$

2.2.4 Strain energy

Since the constitutive relation of the poroelastic medium is linear, the strain energy V associated with the stress part \mathbf{Q}_s of the state vector is just half the sum of the products of the stresses with their corresponding strains,

$$V = \frac{1}{2} \left(\sum_{I=1}^6 \tau_I e_I + p \zeta \right) = \frac{1}{2} (\boldsymbol{\tau}^T \mathbf{e} + p \zeta). \quad (2.81)$$

Using equation (7.132) from [18], $\tau_I = c_{IJE} e_J - \alpha_I p$, we can write $c_{IJE} e_J = \tau_I + \alpha_I p$. (Note that in the principal axes of an orthotropic material, $\alpha_I = 0$ for $I > 3$ — there is no equivalent shear stress associated with the fluid pressure in the principal axes.) Letting \mathbf{S} be the compliance matrix of the drained skeleton — the inverse of the matrix formed by the drained elastic stiffness parameters c_{IJ} — in matrix notation we have

$$\mathbf{e} = \mathbf{S}(\boldsymbol{\tau} + p\boldsymbol{\alpha}). \quad (2.82)$$

Here $\boldsymbol{\tau}$ and \mathbf{e} are arranged as 6×1 column vectors, not as 3×3 matrices.

To get the variation of fluid content ζ , let us return to equation (7.131) of [18], $p = M(\zeta - \alpha^T \mathbf{e})$. In matrix notation, this is $p = M(\zeta - \boldsymbol{\alpha}^T \mathbf{e}) = M(\zeta - \boldsymbol{\alpha}^T \mathbf{S} \boldsymbol{\tau} - p \boldsymbol{\alpha}^T \mathbf{S} \boldsymbol{\alpha})$. Solving for ζ in terms of the stress variables results in

$$\zeta = \left(\frac{1}{M} + \boldsymbol{\alpha}^T \mathbf{S} \boldsymbol{\alpha} \right) p + \boldsymbol{\alpha}^T \mathbf{S} \boldsymbol{\tau}. \quad (2.83)$$

Substituting (2.82) and (2.83) into (2.81), we get

$$V = \frac{1}{2} \left(\boldsymbol{\tau}^T \mathbf{S} \boldsymbol{\tau} + 2p \boldsymbol{\alpha}^T \mathbf{S} \boldsymbol{\tau} + p^2 \left(\frac{1}{M} + \boldsymbol{\alpha}^T \mathbf{S} \boldsymbol{\alpha} \right) \right). \quad (2.84)$$

In matrix form this is

$$V = \frac{1}{2} \mathbf{Q}_s^T \mathbf{E}_s \mathbf{Q}_s, \quad (2.85)$$

where the matrix \mathbf{E}_s is

$$\mathbf{E}_s = \begin{pmatrix} \mathbf{S} & \mathbf{S} \boldsymbol{\alpha} \\ \boldsymbol{\alpha}^T \mathbf{S} & \frac{1}{M} + \boldsymbol{\alpha}^T \mathbf{S} \boldsymbol{\alpha} \end{pmatrix}. \quad (2.86)$$

2.2.5 Kinetic energy

The derivation of kinetic energy from the two-dimensional case carries over directly to three dimensions, and the kinetic energy is

$$T = \frac{1}{2} \mathbf{Q}_v^T \mathbf{E}_v \mathbf{Q}_v, \quad (2.87)$$

where the matrix \mathbf{E}_v is

$$\mathbf{E}_v = \begin{pmatrix} \rho & 0 & 0 & \rho_f & 0 & 0 \\ 0 & \rho & 0 & 0 & \rho_f & 0 \\ 0 & 0 & \rho & 0 & 0 & \rho_f \\ \rho_f & 0 & 0 & m_1 & 0 & 0 \\ 0 & \rho_f & 0 & 0 & m_2 & 0 \\ 0 & 0 & \rho_f & 0 & 0 & m_3 \end{pmatrix}. \quad (2.88)$$

2.2.6 Total energy

Combining (2.85) and (2.87) gives the total energy per unit volume in terms of the state vector \mathbf{Q} as

$$\mathcal{E} = T + V = \frac{1}{2} \mathbf{Q}^T \mathbf{E} \mathbf{Q}, \quad (2.89)$$

where \mathbf{E} is

$$\mathbf{E} = \begin{pmatrix} \mathbf{E}_s & 0_{7 \times 6} \\ 0_{6 \times 7} & \mathbf{E}_v \end{pmatrix}. \quad (2.90)$$

2.2.7 Symmetrization

Using the block structure of the system, $\mathbf{E}\check{\mathbf{A}}$ is symmetric if and only if

$$\mathbf{E}_s\check{\mathbf{A}}_{sv} = (\mathbf{E}_v\check{\mathbf{A}}_{vs})^T. \quad (2.91)$$

After substantial algebra, it can in fact be shown that

$$\mathbf{E}_s\check{\mathbf{A}}_{sv} = - \begin{pmatrix} n_1 & 0 & 0 & 0 & 0 & 0 \\ 0 & n_2 & 0 & 0 & 0 & 0 \\ 0 & 0 & n_3 & 0 & 0 & 0 \\ 0 & n_3 & n_2 & 0 & 0 & 0 \\ n_3 & 0 & n_1 & 0 & 0 & 0 \\ n_2 & n_1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -n_1 & -n_2 & -n_3 \end{pmatrix} \quad (2.92)$$

and

$$\mathbf{E}_v\check{\mathbf{A}}_{vs} = - \begin{pmatrix} n_1 & 0 & 0 & 0 & n_3 & n_2 & 0 \\ 0 & n_2 & 0 & n_3 & 0 & n_1 & 0 \\ 0 & 0 & n_3 & n_2 & n_1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & -n_1 \\ 0 & 0 & 0 & 0 & 0 & 0 & -n_2 \\ 0 & 0 & 0 & 0 & 0 & 0 & -n_3 \end{pmatrix}. \quad (2.93)$$

Thus \mathbf{E} does indeed symmetrize the system.

Since \mathbf{E} symmetrizes the system, we immediately know that the governing equations of three-dimensional poroelasticity are hyperbolic by the argument of Section 2.1.7. An energy norm and energy inner product can also be defined in exactly the same fashion as for two dimensions. Furthermore, we can easily see that $\mathbf{E}\mathbf{D}$ is a symmetric negative-definite matrix in three dimensions as well, since

$$\mathbf{E}\mathbf{D} = \begin{pmatrix} 0_{7 \times 7} & 0_{7 \times 6} \\ 0_{6 \times 7} & \mathbf{E}_v\mathbf{D}_v \end{pmatrix}, \quad (2.94)$$

and

$$\mathbf{E}_v\mathbf{D}_v = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -\eta/\kappa_1 & 0 & 0 \\ 0 & 0 & 0 & 0 & -\eta/\kappa_2 & 0 \\ 0 & 0 & 0 & 0 & 0 & -\eta/\kappa_3 \end{pmatrix}. \quad (2.95)$$

This means that the energy density \mathcal{E} is a strictly convex entropy function in three dimensions as well.

2.3 Linear acoustics

The partial differential equations governing linear acoustics are well-known, and will not be re-derived here. This section will, however, state how they are incorporated into the same framework as poroelasticity.

For this work the PDEs of acoustics will be cast in the same form as the poroelastic system (2.43), with the same state vector; however, in a fluid the variables $\boldsymbol{\tau}$ and \mathbf{v} will be defined to be identically zero. The state variable p will be used for the fluid pressure, and \mathbf{q} for its velocity. (In fact the total stress tensor in the fluid is $-p\mathbf{I}$, but it is more convenient to use the single pressure variable and ignore $\boldsymbol{\tau}$ in the fluid.) The appropriate coefficient matrices have the same block form as for poroelasticity, with blocks for the two-dimensional system given by

$$\mathbf{A}_{sv} = \begin{pmatrix} 0 & 0 & K_f & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \quad \mathbf{A}_{vs} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ \frac{1}{\rho_f} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \quad (2.96)$$

$$\mathbf{B}_{sv} = \begin{pmatrix} 0 & 0 & 0 & K_f \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix} \quad \mathbf{B}_{vs} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ \frac{1}{\rho_f} & 0 & 0 & 0 \end{pmatrix}. \quad (2.97)$$

Blocks for acoustics in the framework of the three-dimensional system (2.79) are given in compact form by

$$\check{\mathbf{A}}_{sv} = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & n_x K_f & n_y K_f & n_z K_f \end{pmatrix}, \quad \check{\mathbf{A}}_{vs} = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \frac{n_x}{\rho_f} \\ 0 & 0 & 0 & 0 & 0 & 0 & \frac{n_y}{\rho_f} \\ 0 & 0 & 0 & 0 & 0 & 0 & \frac{n_z}{\rho_f} \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}. \quad (2.98)$$

The properties are taken to be isotropic, so $\check{\mathbf{A}} = n_x \mathbf{A} + n_y \mathbf{B} + n_z \mathbf{C}$ here is written in terms of a vector $\mathbf{n} = (n_x, n_y, n_z)$ in the global problem coordinates. No source terms are modeled for the fluid; in particular, the dissipation matrix \mathbf{D} is identically zero.

Similarly to a poroelastic medium, it is also possible to write a matrix \mathbf{E} such that the energy density in the fluid is $\mathcal{E} = \frac{1}{2} \mathbf{Q}^T \mathbf{E} \mathbf{Q}$. This \mathbf{E} matrix has the same block structure as for poroelasticity; its blocks in two dimensions are

$$\mathbf{E}_s = \begin{pmatrix} \frac{1}{K_f} & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \end{pmatrix}, \quad \mathbf{E}_v = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & 0 & \rho_f & 0 \\ 0 & 0 & 0 & \rho_f \end{pmatrix}, \quad (2.99)$$

and in three dimensions,

$$\mathbf{E}_s = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \frac{1}{K_f} \end{pmatrix}, \quad \mathbf{E}_v = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & \rho_f & 0 & 0 \\ 0 & 0 & 0 & 0 & \rho_f & 0 \\ 0 & 0 & 0 & 0 & 0 & \rho_f \end{pmatrix}. \quad (2.100)$$

Note that the \mathbf{E} defined this way is only positive-semidefinite, not positive-definite as for poroelasticity. However, the null space of \mathbf{E} consists only of the variables that are defined to be identically zero in the fluid, so it is essentially positive-definite, and it is still sensible to use it to define an energy inner product and norm in the fluid. Just as for poroelasticity, \mathbf{E} symmetrizes the first-order hyperbolic systems for acoustics in both two and three dimensions.

2.4 Interface conditions

Nontrivial conditions relating the state variables on either side of an interface between distinct poroelastic materials have been proposed by a number of authors. These include Deresiewicz and Skalak [31], who proposed an imperfect hydraulic contact condition relating the pressure difference across the interface to the normal fluid flow rate, and Sharma [69], who formulated a loose contact condition modeling lubricated slippage between the two sides. This work uses Deresiewicz and Skalak's imperfect hydraulic contact condition; Deresiewicz and Skalak showed that this condition is sufficient to give a unique solution to Biot's equations in a discontinuous medium, and Gurevich and Schoenberg [43] examined how such a condition could arise asymptotically from a smoothly varying medium as the region over which the material properties vary is shrunk to zero thickness.

The imperfect hydraulic contact condition can be written in the form

$$\begin{aligned} \boldsymbol{\tau}_l \cdot \mathbf{n} &= \boldsymbol{\tau}_r \cdot \mathbf{n} \\ \mathbf{v}_l &= \mathbf{v}_r \\ \mathbf{q}_l \cdot \mathbf{n} &= \mathbf{q}_r \cdot \mathbf{n} \\ p_l - p_r &= \frac{1}{\mathcal{K}} \widehat{\mathbf{q}} \cdot \mathbf{n}. \end{aligned} \quad (2.101)$$

Here the subscripts l and r represent the left and right sides of the interface, chosen arbitrarily, and \mathbf{n} is the unit normal to the interface, pointing from left to right. These equations have direct physical significance: the first is a statement of the continuity of traction across the interface, the second states that the materials stay connected to each other, the third requires that all fluid entering the interface should exit the other side, and the fourth relates the fluid flow rate across the interface to the pressure difference forcing it across. The parameter \mathcal{K} is a measure of the ability of the interface to conduct fluid, and ranges from zero, representing a completely impermeable interface, to $+\infty$, representing no impedance

to fluid flow. The quantity $\widehat{\mathbf{q}} \cdot \mathbf{n}$ is the volume flow rate of fluid across the interface, which is equal to both $\mathbf{q}_l \cdot \mathbf{n}$ and $\mathbf{q}_r \cdot \mathbf{n}$ according to the third equation of (2.101). This ambiguity will be revisited in the discussion of the implementation of this interface condition.

Since both $\mathcal{K} = 0$ and $\mathcal{K} = +\infty$ are common and important cases, and since infinite values are inconvenient in numerical computations, it is useful to reparameterize the fourth equation of (2.101). Noting that $1/\mathcal{K}$ has the same units as acoustic impedance, define

$$\frac{1}{\mathcal{K}} =: Z_f \frac{1 - \eta_d}{\eta_d}, \quad (2.102)$$

where Z_f is the impedance of the pore fluid in the left medium and $\eta_d \in [0, 1]$ is a new nondimensional parameter termed the *interface discharge efficiency*. Setting $\mathcal{K} = +\infty$ now corresponds to setting $\eta_d = 1$, and $\mathcal{K} = 0$ corresponds to $\eta_d = 0$. Substituting this into the fourth equation of (2.101) and multiplying through by η_d results in

$$\begin{aligned} \boldsymbol{\tau}_l \cdot \mathbf{n} &= \boldsymbol{\tau}_r \cdot \mathbf{n} \\ \mathbf{v}_l &= \mathbf{v}_r \\ \mathbf{q}_l \cdot \mathbf{n} &= \mathbf{q}_r \cdot \mathbf{n} \\ \eta_d(p_l - p_r) &= Z_f(1 - \eta_d)\widehat{\mathbf{q}} \cdot \mathbf{n}, \end{aligned} \quad (2.103)$$

which presents no special difficulty for any value of η_d , including 0 and 1. This is the form of the interface condition that will be used in numerical computations.

Between a poroelastic medium and a fluid, a similar condition can be used to model an interface that may resist the flow of fluid; this condition can be written as

$$\begin{aligned} \mathbf{q}_f \cdot \mathbf{n} &= (\mathbf{v}_p + \mathbf{q}_p) \cdot \mathbf{n} \\ -p_f \mathbf{n} &= \boldsymbol{\tau}_p \cdot \mathbf{n} \\ p_p - p_f &= \frac{1}{\mathcal{K}} \mathbf{q}_p \cdot \mathbf{n}. \end{aligned} \quad (2.104)$$

Here, the subscript f indicates quantities in the fluid, while p indicates quantities in the poroelastic medium. The vector \mathbf{q}_f represents the fluid velocity, and the unit interface normal \mathbf{n} points from the poroelastic medium into the fluid. These equations have meanings similar to the poroelastic interface equations (2.101): the first states continuity of fluid flow, the second, continuity of traction, and the third relates fluid flow rate to pressure difference. The parameter \mathcal{K} has the same meaning as before, though its value may be different. This condition seems to have been first used at a poroelastic-fluid interface by Rosenbaum [67], and has since been employed by other authors [10, 36]. Chiavassa and Lombard [22] in particular use this method in their numerical work, and also demonstrate that the resulting coupled systems of PDEs on the fluid and poroelastic domains are well-posed. For easier implementation, it is again useful to replace $1/\mathcal{K}$ with $Z_f(1 - \eta_d)/\eta_d$, where Z_f is the acoustic impedance of the fluid medium, giving the alternative form

$$\begin{aligned} \mathbf{q}_f \cdot \mathbf{n} &= (\mathbf{v}_p + \mathbf{q}_p) \cdot \mathbf{n} \\ -p_f \mathbf{n} &= \boldsymbol{\tau}_p \cdot \mathbf{n} \\ \eta_d(p_p - p_f) &= Z_f(1 - \eta_d)\mathbf{q}_p \cdot \mathbf{n}. \end{aligned} \quad (2.105)$$

The reader should note that these conditions are expressed independently of the number of spatial dimensions — they are equally valid in two-dimensional and three-dimensional systems.

2.5 Decay of poroelastic waves in practical problems

One potential issue with poroelasticity as a practical theory is that, while the slow P wave is distinctive and adds new behavior compared to ordinary elastodynamics, it is rapidly damped, leaving behind only the fast P and S waves. The remaining waves, on the other hand, are lightly damped, and at least over short distances it may not be possible to realistically distinguish them from waves in a nonporous elastic medium. The aim of this section is to investigate the decay rates of poroelastic waves in the context of seismic waves and bone modeling, in order to determine how important the slow P wave and the dissipation of the other wave families will be for these types of problems, and to thereby evaluate whether poroelasticity theory is useful in these contexts.

The quantity used here to evaluate decay is the *characteristic decay length* L_d of a time-harmonic wave train, defined as the distance over which the amplitude of the waves decreases by a factor of e . The plane wave solutions of Section 4.1 are used to determine this decay length — for these solutions, it can easily be computed as $L_d = 1/|\text{Im } k|$, where k is the complex wavenumber of the plane wave.

2.5.1 Seismic waves

For seismic wave problems, the orthotropic sandstone and isotropic shale of Table A.1 are used as representative materials. The isotropic sandstone was deemed redundant and omitted. Figure 2.1 shows the decay length as a function of frequency for the fast P and S waves in these materials, while Figure 2.2 shows it for the slow P wave; the frequency ranges for all plots are from 1 Hz up to the cutoff frequency for validity of low-frequency Biot theory. Contour plots are used to display the results for sandstone because it is anisotropic, so its characteristic decay lengths depend on propagation direction as well as frequency.

In a seismic setting, wave frequencies will typically be at the low end of the range of Figure 2.1, up to a few hertz or the low tens of hertz, while propagation distances will be in the hundreds of meters to multiple kilometers. At this combination of length scale and frequency, the decay of the fast P and S waves due to viscosity of the pore fluid should be negligible. Meanwhile, Figure 2.2 shows that the slow P wave will not play a directly observable role in such problems, although it may still have an indirect effect if it is created at an inhomogeneity — for instance at the boundary between a porous rock formation and a non-porous one — and carries away energy that would otherwise have gone into the other, longer-lived wave families. For stone media saturated with brine, then, poroelastic effects are not significant, except inasmuch as they may change reflection and transmission coefficients at material boundaries. Note that this evaluation does not consider oil-bearing rock; depending on the temperature and type of oil, crude oil may have a viscosity 4-2000 times as great as water [1], so damping of fast P and S waves may be noticeable at higher frequencies and at the higher end of this viscosity range.




Figure 2.1: Characteristic decay lengths for fast P and S waves in the stone media of Table A.1. For the orthotropic sandstone, propagation direction is measured as the angle of the wavevector counterclockwise from the material principal 1 axis.




Figure 2.2: Characteristic decay lengths for slow P waves in the stone media of Table A.1. For the orthotropic sandstone, propagation direction is measured as the angle of the wavevector counterclockwise from the material principal 1 axis.

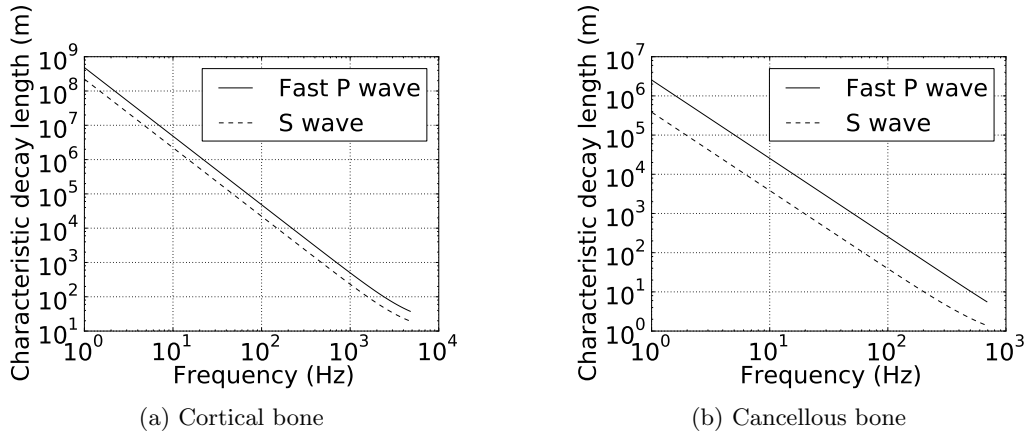


Figure 2.3: Characteristic decay lengths for fast P and S waves in the bone media of Table A.1.




Figure 2.4: Characteristic decay lengths for fast P and S waves in the bone media of Table A.1.

2.5.2 Bone problems

The bone media assessed here are the cortical and cancellous bone of Table A.1. Note that the cortical bone properties are for the vascular pore space. Figure 2.3 shows the decay lengths of the fast P and S waves, with Figure 2.4 covering the slow P wave. The frequency range is again from 1 Hz to the low-frequency validity cutoff.

For bone, problem length scales will be in the millimeters, up to perhaps a few centime-

ters, and the relevant frequencies will be high, at least in the kilohertz range. (In fact, to truly perform this assessment for ultrasound or shock waves, it would likely be necessary to use a model valid at much higher frequencies, but such a model is beyond the scope of this work.) For these length scales and frequencies, Figure 2.3 shows that the fast P and S waves will experience very little attenuation, but from Figure 2.4 we can see that the slow P wave will also not be heavily attenuated — even at the highest frequencies in the plots, the slow P wave can easily penetrate centimeters of bone. This suggests that using the full poroelasticity system can be important in modeling the propagation of waves in bone, and that all three types of poroelastic wave could become involved in the response of bone to external excitation.

Chapter 3

METHODS FOR POROELASTICITY ON TWO-DIMENSIONAL
RECTILINEAR GRIDS

Now that the governing equations of poroelastic-fluid systems have been laid down and some of their properties examined, it is time to begin formulating a numerical method to solve them. This chapter describes the numerical methods and software used for the initial work simulating poroelastic wave propagation on rectilinear grids. It corresponds to section 3 of the author’s paper with Ou and LeVeque [56].

3.1 Wave propagation

The work of this chapter and Chapter 4 solves the equations of poroelasticity using a rectilinear grid finite volume approach. This chapter describes the basics of the finite volume method used on rectilinear grids, as well as specifics of how it is applied to poroelasticity. For a comprehensive discussion of this class of finite volume methods, see LeVeque’s book [57].

The class of finite volume method used here updates cell averages at every step by solving a Riemann problem — a Cauchy problem with piecewise constant initial data — between each pair of adjacent grid cells. Thinking of one cell as the “left” cell of the problem, and the other as the “right” cell, the Riemann solution process produces the left-going and right-going fluctuations $\mathcal{A}^- \Delta \mathbf{Q}$ and $\mathcal{A}^+ \Delta \mathbf{Q}$ — the rates of change of the cell variables \mathbf{Q} caused by the left-going and right-going waves — along with a set of waves \mathcal{W}_i with speeds s_i that are used to implement higher-order correction terms. With these correction terms included, the methods used here are second-order accurate. Where solutions are not smooth, wave limiters can be used on the higher-order terms to prevent spurious oscillations. While limiters can reduce the asymptotic order of accuracy of the solution, they often decrease the actual value of the error, depending on the norm being used to measure it and on the grid resolution. This means that they are not desirable for convergence studies, because they can obscure the behavior of the other components of the numerical method, but they are very useful for practical problems. They can also improve the qualitative behavior of the solution by suppressing dispersive errors, leading to improved estimates of quantities such as wave arrival times, and keeping total variation from increasing.

For a homogeneous first-order hyperbolic system such as (2.57), the left-going and right-going fluctuations are related to the waves and wave speeds by

$$\mathcal{A}^+ \Delta \mathbf{Q} = \sum_{s_i > 0} s_i \mathcal{W}_i, \quad \mathcal{A}^- \Delta \mathbf{Q} = \sum_{s_i < 0} s_i \mathcal{W}_i. \quad (3.1)$$

For linear problems, including all the hyperbolic systems in the poroelastic-fluid models addressed here, the waves are simply eigenvectors of the flux Jacobian matrix (for instance, the matrix \mathbf{A} of (2.45) for waves propagating in the x direction in a material with its

principal axes aligned with the coordinate axes) associated with the material through which the wave propagates — that is, they have the form $\mathcal{W}_i = \beta_i \mathbf{r}_i$, where \mathbf{r}_i is the eigenvector and β_i is a scalar that gives the strength of the wave. Each wave speed s_i is the corresponding eigenvalue of the flux Jacobian.

A quantity of critical importance in these solution methods is the *CFL number* ν . Informally, the CFL number is the ratio of the distance a wave travels in one time step to the width of a grid cell; more formally, for a rectilinear grid the global CFL number is

$$\nu = \max_{\text{all cells, all waves}} \max \left(\frac{|s_x| \Delta t}{\Delta x}, \frac{|s_z| \Delta t}{\Delta z} \right). \quad (3.2)$$

Here s_x and s_z are the speeds of waves generated from the Riemann problems in the x and z directions, Δx and Δz are the grid spacings, and Δt is the time step size. The methods used here are stable for $\nu \leq 1$; since ν comes from a maximum over all waves, this means that the stability of this method in a poroelastic medium is limited by the fast P wave.

Because the poroelasticity equations are a linear system, solution of the Riemann problem is straightforward. There is one complication, however. The scope of the work here includes domains composed of multiple materials — in fact, the simulation code implementing these methods is designed with the capability for each grid cell to be made of a different material — so the coefficient matrices are only piecewise constant. The cell boundaries are always chosen to coincide with the material boundaries, but Riemann problems must still be solved between domains with different coefficient matrices \mathbf{A} and \mathbf{B} . Consider a Riemann problem in the x -direction, with $\mathbf{A} = \mathbf{A}_l$ in the left cell and $\mathbf{A} = \mathbf{A}_r$ in the right cell. In the left cell, the Riemann solution consists of waves with strength β_{li} in the directions of the eigenvectors \mathbf{r}_{li} of \mathbf{A}_l , corresponding to the negative eigenvalues of \mathbf{A}_l ; similarly, the right cell will have waves with strength β_{ri} in the directions of eigenvectors \mathbf{r}_{ri} or \mathbf{A}_r , corresponding to the positive eigenvalues of \mathbf{A}_r . There will also be a stationary discontinuity at the cell interface, which will lie in the null space of \mathbf{A}_l and \mathbf{A}_r . (The fact that \mathbf{A} has the same null space for any poroelastic material greatly simplifies matters here, and the corresponding eigenvectors \mathbf{r}_4 and \mathbf{r}_5 will not carry a subscript identifying them with the left or right material.) The total jump in \mathbf{Q} across all the waves and the stationary discontinuity must add up to the difference in \mathbf{Q} between the left and right states, $\Delta \mathbf{Q} = \mathbf{Q}_r - \mathbf{Q}_l$, so the following equation holds:

$$\Delta \mathbf{Q} = \sum_{i=1}^3 \beta_{li} \mathbf{r}_{li} + \sum_{i=4}^5 \beta_i \mathbf{r}_i + \sum_{i=6}^8 \beta_{ri} \mathbf{r}_{ri} =: \tilde{\mathbf{R}} \boldsymbol{\beta}. \quad (3.3)$$

The wave strengths can thus be computed as $\boldsymbol{\beta} = \tilde{\mathbf{R}}^{-1} \Delta \mathbf{Q}$. (Note that in this formula, the eigenvectors are ordered by eigenvalues, from most negative to most positive.) In practice, since the strength of the stationary discontinuity is never used directly, β_4 and β_5 are never computed. The same analysis holds for a Riemann problem in the z -direction. This approach corresponds to an open-pore condition between the two poroelastic media, as described by Deresiewicz and Skalak [31] and validated by Gurevich and Schoenberg [43]. Other interface conditions will be incorporated into the formulation in Chapter 5 and put into action in chapters 6 and 7.

Because the eigenstructure of poroelasticity is somewhat complex, the early version of the simulation code used for rectilinear grid results does not compute the eigensystems of \mathbf{A} and \mathbf{B} analytically. Instead, it uses LAPACK [3] to compute the eigenvalues and eigenvectors of \mathbf{A} and \mathbf{B} for every poroelastic material present in the model, and the $\tilde{\mathbf{R}}^{-1}$ matrices for each Riemann solve direction and every pair of left and right materials that could occur. For efficiency, it pre-computes these quantities for all materials used (or all possible pairs of materials in the case of $\tilde{\mathbf{R}}^{-1}$) before starting the solution proper, and looks them up using a material number stored with each cell during the Riemann solves.

Boundary conditions are implemented using the usual ghost cell approach [57]. Ghost cell values are set using either zero-order extrapolation, for boundaries where waves should flow outward and not return, or by setting the ghost cell values equal to the exact solution at the centers of those cells, when the code is verified against known analytic solutions.

3.2 Operator splitting

The dissipative part $\mathbf{Q}_t = \mathbf{D}\mathbf{Q}$ of the poroelasticity equations is included using operator splitting. Since the \mathbf{D} matrix is constant, the exact solution operator $\exp(\mathbf{D}\Delta t)$ can be used to advance the solution by a time increment Δt ; not only is this the most accurate solution available for this part of the system, it is also unconditionally stable and allows the time step to be chosen based solely on stability for the wave propagation part of the system.

The software framework used offers either Godunov or Strang splitting as a run-time option. Godunov splitting is formally first-order accurate in time and uses a single full-length step of the source term operator per time step, while Strang splitting is second-order and uses two half-steps of the source term. For many practical problems Godunov splitting is a good choice because it displays similar error to Strang — the coefficient of the first-order error term is often small — while being less computationally intensive. However, Strang splitting is primarily used here because it displays substantially greater accuracy for the particular poroelasticity problems solved, and because the source term is computationally cheap compared to the wave propagation part of the system. For comparison, results are also shown for Godunov splitting.

3.3 Stiff regime and subcharacteristic condition

For some cases considered in this work, the time step is much larger than the characteristic time scales associated with the solution of $\partial_t \mathbf{Q} = \mathbf{D}\mathbf{Q}$. These cases fall outside the regime where asymptotic error estimates for small time steps are relevant, and for them the source term is stiff, a known source of difficulty in operator splitting approaches for hyperbolic equations [24, 58]. Based on a conjecture of Pember [66], one can expect to avoid spurious solutions for this stiff relaxation system if the poroelasticity equations (2.43) satisfy a *subcharacteristic condition*, where the wave speeds for the reduced equations obtained by restricting the full system to the equilibrium manifold of the dissipation term (i.e. zero fluid velocity relative to the solid matrix) interlace with the wave speeds for the full system. The appropriate generalization of Pember’s conjecture to systems of more than two equations is not obvious, but based on the principle that information should propagate more slowly (cer-

tainly no more quickly!) in the reduced system than in the full system, spurious solutions should be prevented if for all possible wave propagation directions the speeds λ_1 and λ_2 of the reduced system are strictly less than the speed of the fastest wave of the full system,

$$\lambda_1 < c_{pf}, \quad \lambda_2 < c_{pf}. \quad (3.4)$$

Here c_{pf} is the speed of a fast P wave. Negative eigenvalues are ignored for both the full and reduced systems, because they are simply the negatives of the wave speeds and will automatically satisfy a similar inequality. Zero eigenvalues of the full system are also ignored, since they correspond to eigencomponents of the solution that are left unchanged in the wave propagation part of the solution process, and are evolved according to the exact solution operator $\exp(\mathbf{D}\Delta t)$ in the dissipation part.

The appropriate reduced system can be constructed following the derivation of Chen, Levermore, and Liu [21]. First, examine the dissipation part of the system in isolation,

$$\partial_t \mathbf{Q} = \mathbf{D}\mathbf{Q}. \quad (3.5)$$

System (3.5) has six conserved quantities $\mathbf{u} = \mathbf{\Pi}\mathbf{Q}$, which are related to the state variables \mathbf{Q} by the matrix

$$\mathbf{\Pi} := \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & \rho_f/\rho & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & \rho_f/\rho \end{pmatrix}. \quad (3.6)$$

The fact that \mathbf{u} is a vector of conserved quantities of (3.5) follows immediately from the fact that $\mathbf{\Pi}\mathbf{D} = 0$, so that $\partial_t \mathbf{u} = \mathbf{\Pi}\partial_t \mathbf{Q} = \mathbf{\Pi}\mathbf{D}\mathbf{Q} = 0$. Given the conserved quantities \mathbf{u} , the unique equilibrium \mathbf{Q}_{eq} of (3.5) that satisfies $\mathbf{D}\mathbf{Q}_{\text{eq}} = 0$ and $\mathbf{\Pi}\mathbf{Q}_{\text{eq}} = \mathbf{u}$ is $\mathbf{Q}_{\text{eq}} = \mathbf{G}\mathbf{u}$, where

$$\mathbf{G} := \begin{pmatrix} \mathbf{I}_{6 \times 6} \\ \mathbf{0}_{2 \times 6} \end{pmatrix}. \quad (3.7)$$

Notice that $\mathbf{\Pi}$ and \mathbf{G} satisfy the relation $\mathbf{\Pi}\mathbf{G} = \mathbf{I}_{6 \times 6}$. The reduced system is found by multiplying the full poroelastic system from the left by $\mathbf{\Pi}$ to eliminate the dissipation term, and requiring the state vector \mathbf{Q} to lie on the equilibrium manifold, $\mathbf{Q} = \mathbf{Q}_{\text{eq}} = \mathbf{G}\mathbf{u}$, resulting in

$$\partial_t \mathbf{u} + \mathbf{\Pi}\mathbf{A}\mathbf{G}\partial_x \mathbf{u} + \mathbf{\Pi}\mathbf{B}\mathbf{G}\partial_z \mathbf{u} = 0. \quad (3.8)$$

With the matrix $\mathbf{\Pi}$ in hand, it is now possible to show that the additional condition of Chen, Levermore, and Liu mentioned in Section 2.1.8 also holds — that the statements $\mathbf{D}\mathbf{Q} = 0$ and $\Phi'(\mathbf{Q})^T \mathbf{D}\mathbf{Q} = \mathbf{Q}^T \mathbf{E}\mathbf{D}\mathbf{Q} = 0$ are equivalent to $\Phi'(\mathbf{Q})^T = \mathbf{Q}^T \mathbf{E} = \mathbf{v}^T \mathbf{\Pi}$ for some $\mathbf{v} \in \mathbb{R}^6$. First, if $\mathbf{Q}^T \mathbf{E} = \mathbf{v}^T \mathbf{\Pi}$, we immediately have $\mathbf{Q}^T \mathbf{E}\mathbf{D}\mathbf{Q} = \mathbf{v}^T \mathbf{\Pi}\mathbf{D}\mathbf{Q} = 0$ since $\mathbf{\Pi}\mathbf{D} = 0$. Section 2.1.8 also demonstrated that if $\mathbf{D}\mathbf{Q} = 0$, then $q_x = q_z = 0$. Thus the fifth, sixth, seventh, and eighth components of $\mathbf{Q}^T \mathbf{E}$ are $(\mathbf{Q}^T \mathbf{E})_5 = \rho v_x$, $(\mathbf{Q}^T \mathbf{E})_6 = \rho v_z$, $(\mathbf{Q}^T \mathbf{E})_7 = \rho_f v_x$, and $(\mathbf{Q}^T \mathbf{E})_8 = \rho_f v_z$. We therefore have $\mathbf{Q}^T \mathbf{E} = \mathbf{v}^T \mathbf{\Pi}$ with \mathbf{v} given by

$$\mathbf{v} = ((\mathbf{Q}^T \mathbf{E})_1 \quad (\mathbf{Q}^T \mathbf{E})_2 \quad (\mathbf{Q}^T \mathbf{E})_3 \quad (\mathbf{Q}^T \mathbf{E})_4 \quad \rho v_x \quad \rho v_z)^T. \quad (3.9)$$

Given that a strictly convex entropy function $\Phi(\mathbf{Q}) = \mathcal{E}(\mathbf{Q})$ exists and the above condition is satisfied, Theorem 2.1 and the subsequent remark of Chen, Levermore, and Liu [21] imply that the reduced system (3.8) is hyperbolic, and satisfies a nonstrict subcharacteristic condition, which we can render here in the context of poroelasticity as

$$c_{ps} \leq \lambda_2 \leq c_{pf}, \quad 0 \leq \lambda_1 \leq c_s. \quad (3.10)$$

While this does not imply the strict subcharacteristic condition (3.4), it is nearly as useful. This characterization of the eigenvalues comes from the fact that eigenvalues of the symmetric generalized eigenproblem (2.59) satisfy the Rayleigh quotient minimax principle

$$\lambda_{fk}^\downarrow = \max_{\mathcal{S}_k} \min_{\mathbf{v} \in \mathcal{S}_k} \frac{\mathbf{v}^T \mathbf{E} \check{\mathbf{A}} \mathbf{v}}{\mathbf{v}^T \mathbf{E} \mathbf{v}}, \quad \lambda_{fk}^\uparrow = \min_{\mathcal{S}_k} \max_{\mathbf{v} \in \mathcal{S}_k} \frac{\mathbf{v}^T \mathbf{E} \check{\mathbf{A}} \mathbf{v}}{\mathbf{v}^T \mathbf{E} \mathbf{v}}, \quad (3.11)$$

where \mathcal{S}_k is any k -dimensional subspace of \mathbb{R}^8 , λ_{fk}^\downarrow is the k 'th eigenvalue of the full system counting down from the largest, and λ_{fk}^\uparrow is the k 'th eigenvalue counting up from the smallest. Chen, Levermore, and Liu prove that the eigenvalues of the reduced system satisfy a similar minimax principle over a restricted set of subspaces,

$$\lambda_{ek}^\downarrow = \max_{\mathcal{S}_k \subseteq \mathcal{R}(\mathbf{G})} \min_{\mathbf{v} \in \mathcal{S}_k} \frac{\mathbf{v}^T \mathbf{E} \check{\mathbf{A}} \mathbf{v}}{\mathbf{v}^T \mathbf{E} \mathbf{v}}, \quad \lambda_{ek}^\uparrow = \min_{\mathcal{S}_k \subseteq \mathcal{R}(\mathbf{G})} \max_{\mathbf{v} \in \mathcal{S}_k} \frac{\mathbf{v}^T \mathbf{E} \check{\mathbf{A}} \mathbf{v}}{\mathbf{v}^T \mathbf{E} \mathbf{v}}, \quad (3.12)$$

where the subscript e indicates eigenvalues of the reduced (“equilibrium”) system and $\mathcal{R}(\mathbf{G})$ is the range space of \mathbf{G} .

Equality can be realized in the nonstrict subcharacteristic condition (3.10) — for example, if the fluid density ρ_f for the orthotropic sandstone material whose parameters are given in the first column of Table A.1 is reduced from 1040 kg/m³ to 208.9 kg/m³, a fast P wave traveling along the material principal 1-axis shows no fluid relative motion. This means its eigenvector lies in $\mathcal{R}(\mathbf{G})$, so by (3.12) the reduced system shares the same wave speed. Fortunately, this turns out to be innocuous from the standpoint of the true solution — if the eigenvector \mathbf{v} associated with some wave is in $\mathcal{R}(\mathbf{G})$, then it is in $\mathcal{N}(\mathbf{D})$; since it is an eigenvector of both parts of the system, the corresponding eigencomponent of the solution can be decoupled from the rest of the system, and its solution is independent of the other eigencomponents. Furthermore, the PDE describing the evolution of this component of the solution is a purely hyperbolic one, with no source term — there is no momentum transfer due to viscous drag between the solid and fluid for this wave mode because there is no relative motion between them. A similar result carries over to the numerical solution obtained by operator splitting: if \mathbf{v} is in $\mathcal{N}(\mathbf{D})$, then $\exp(\mathbf{D}\Delta t)\mathbf{v} = \mathbf{v}$, and the corresponding eigencomponent of \mathbf{Q} passes through the solution operator for the dissipation term unchanged.

As an aside, the reduced system (3.8) has a familiar form. Multiplying out the coefficient matrices gives

$$\mathbf{\Pi AG} = \begin{pmatrix} 0 & 0 & 0 & 0 & M\alpha_1 & 0 \\ 0 & 0 & 0 & 0 & -c_{11}^u & 0 \\ 0 & 0 & 0 & 0 & -c_{13}^u & 0 \\ 0 & 0 & 0 & 0 & 0 & -c_{55}^u \\ 0 & -\frac{1}{\rho} & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & -\frac{1}{\rho} & 0 & 0 \end{pmatrix}, \quad \mathbf{\Pi BG} = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & M\alpha_3 \\ 0 & 0 & 0 & 0 & 0 & -c_{13}^u \\ 0 & 0 & 0 & 0 & 0 & -c_{33}^u \\ 0 & 0 & 0 & 0 & -c_{55}^u & 0 \\ 0 & 0 & 0 & -\frac{1}{\rho} & 0 & 0 \\ 0 & 0 & -\frac{1}{\rho} & 0 & 0 & 0 \end{pmatrix}. \quad (3.13)$$

The lower-right 5×5 portions of these matrices are just the coefficient matrices for orthotropic plane-strain elasticity, with the fluid pressure coming along as an additional variable determined entirely by the elastic field variables. Because of this the faster wave of the reduced system will be identified as the “reduced P wave” and the slower one as the “reduced S wave.” Note that for this work, the reduced system is only of theoretical importance — for the actual numerical code, it is the full system (2.43) that is discretized.

3.4 Numerical software

The numerical solution techniques described here are implemented using the CLAWPACK finite volume method package, version 4.6 [71]. CLAWPACK implements the parts of a high-resolution finite volume code that are common across all problems, leaving the user to write only problem-specific code such as Riemann solvers. Operator splitting is supported for source terms, such as the dissipative term here, by means of a user-supplied subroutine that advances the system by a specified time step under the action of the source term. Both Godunov and Strang splitting are available. Block-structured Berger-Colella-Oliger adaptive mesh refinement (AMR) is available from the AMRCLAW package [6]; AMRCLAW can also run in parallel on shared-memory systems using OpenMP. Besides rectilinear grids, CLAWPACK and AMRCLAW also support logically rectangular mapped grids, which are used from Chapter 5 onward.

Chapter 4

RESULTS FOR POROELASTICITY ON TWO-DIMENSIONAL RECTILINEAR GRIDS

Results are presented here for four classes of problems. First, convergence of the numerical solution to known analytic plane wave solutions for an orthotropic medium is demonstrated for the wave propagation part of the system alone. The viscous dissipation term is then included, and the effect of operator splitting on accuracy is examined, again comparing against known analytic plane wave solutions. Next, results are shown for simple point sources in uniform orthotropic media, solving test problems previously addressed by de la Puente et al. [29] and Carcione [17] in order to further verify our code. Finally, a larger-scale problem is solved with a domain composed of two isotropic materials, involving wave reflection, refraction, and interconversion at the material boundary, as well as demonstrating the use of adaptive mesh refinement to reduce the time required for solution. This chapter corresponds to Section 4 of the author's paper with Ou and LeVeque [56], and the analytical solution of Section 4.1 is Ou's work.

4.1 Analytical plane wave solution

Before testing the code's convergence against analytic plane wave solutions, however, it is necessary to outline the procedure used to obtain these analytic solutions. Start by assuming the velocity and stress fields have a plane wave form,

$$\mathbf{V} := (v_x \quad v_z \quad q_x \quad q_z)^T = \mathbf{V}_0 \exp(i(k_x x + k_z z - \omega t)) \quad (4.1)$$

$$\mathbf{T} := (\tau_{xx} \quad \tau_{zz} \quad \tau_{xz} \quad -p)^T = \mathbf{T}_0 \exp(i(k_x x + k_z z - \omega t)). \quad (4.2)$$

Here \mathbf{V}_0 and \mathbf{T}_0 are constant vectors, and ω is the prescribed angular frequency of the wave. The wavenumbers k_x and k_z are yet to be determined, but they are related by $k_x = kl_x$ and $k_z = kl_z$, where l_x and l_z are the prescribed, real-valued direction cosines of the wavevector, with $l_x^2 + l_z^2 = 1$.

With these assumptions on the solution, stress-strain equations (2.31) through (2.34) imply

$$-\omega \mathbf{T}_0 = k \mathbf{F} \mathbf{V}_0, \quad (4.3)$$

where the matrix \mathbf{F} is

$$\mathbf{F} = \begin{pmatrix} l_x c_{11}^u & l_z c_{13}^u & \alpha_1 M l_x & \alpha_1 M l_z \\ l_x c_{13}^u & l_z c_{33}^u & \alpha_3 M l_x & \alpha_3 M l_z \\ l_z c_{55}^u & l_x c_{55}^u & 0 & 0 \\ \alpha_1 M l_x & \alpha_3 M l_z & M l_x & M l_z \end{pmatrix}. \quad (4.4)$$

Equations of motion (2.35) through (2.38) also imply

$$k\mathbf{L}\mathbf{T}_0 = -\omega\mathbf{\Gamma}\mathbf{V}_0, \quad (4.5)$$

where the matrices \mathbf{L} and $\mathbf{\Gamma}$ are

$$\mathbf{L} = \begin{pmatrix} l_x & 0 & l_z & 0 \\ 0 & l_z & l_x & 0 \\ 0 & 0 & 0 & l_x \\ 0 & 0 & 0 & l_z \end{pmatrix}, \quad \mathbf{\Gamma} = \begin{pmatrix} \rho & 0 & \rho_f & 0 \\ 0 & \rho & 0 & \rho_f \\ \rho_f & 0 & iY_1(-\omega)/\omega & 0 \\ 0 & \rho_f & 0 & iY_3(-\omega)/\omega \end{pmatrix}, \quad (4.6)$$

and $Y_j(\omega) := i\omega m_j + \eta/\kappa_j$ for $j = 1, 3$.

Combining equations (4.3) and (4.5) yields an eigenproblem for \mathbf{V}_0 and $(\frac{\omega}{k})^2$:

$$\mathbf{\Gamma}^{-1}\mathbf{L}\mathbf{F}\mathbf{V}_0 = \left(\frac{\omega}{k}\right)^2 \mathbf{V}_0. \quad (4.7)$$

Plane wave solutions are obtained by solving this eigenproblem, choosing the eigenvalue $(\frac{\omega}{k})^2$ and eigenvector \mathbf{V}_0 corresponding to the wave family of interest, then backing out the appropriate \mathbf{T}_0 from (4.3). Note that the eigenvalues and eigenvectors will be complex-valued if dissipation is present.

4.2 Plane wave convergence results — inviscid

The plane wave convergence tests are conducted on a uniform domain composed of orthotropic layered sandstone, whose properties are given in Table A.1. The wave-propagation part of the code was first tested alone, without viscous dissipation included. These tests were conducted using plane waves with a fixed angular frequency of 10^4 rad/s, in a square domain 8 m on a side. This distance is about two wavelengths of the fast P wave at this frequency. The total simulation time is $2\pi \times 10^{-4}$ s — one period of the wave. Because there is no intrinsic time scale associated with Biot's equations when viscous dissipation is omitted, the frequency of the plane wave is not directly relevant to the accuracy of these tests; instead, only the number of grid cells per wavelength and the ratio of the simulation time to the wave period are relevant. Boundary conditions are set by filling the ghost cells with the value of the true plane wave solution at the cell centers, and the initial condition on the grid is set the same way. This ensures that the accuracy of the numerical solution is governed only by the correctness of the wave-propagation algorithm used within the problem domain, not by the implementation of the boundary conditions.

Since the code is intended for simulating wave propagation in an orthotropic material, rather than an isotropic one, the speed and associated eigenvector for a plane wave depend on its propagation direction — the solutions for plane waves propagating in different directions are not simply rotated versions of each other, and in order to be confident in the correctness of the code it is necessary to test it with plane waves propagating at a variety of angles θ_{wave} relative to the global x axis. In order to demonstrate correctness when the principal material axes do not coincide with the global coordinate axes, it is also necessary to test the code with a variety of angles θ_{mat} between the material 1 axis and the x axis. Figure

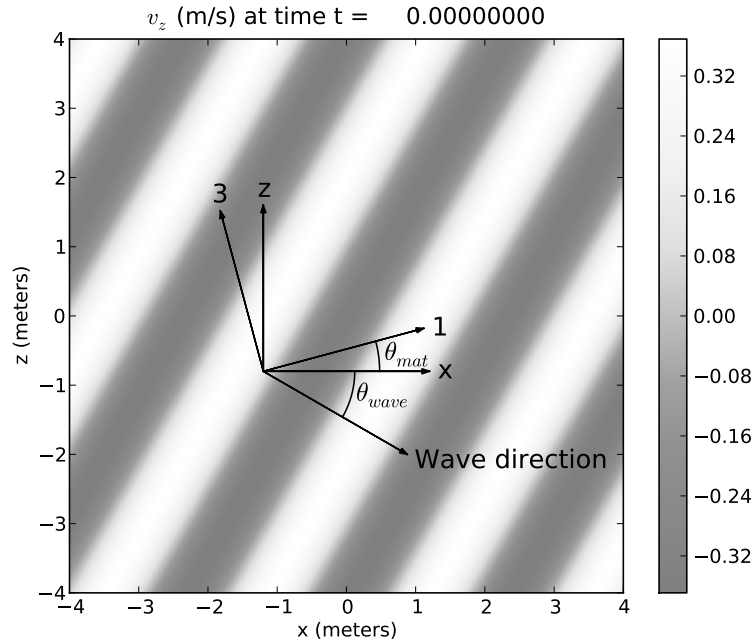


Figure 4.1: Sample plane wave initial condition, showing the relation between the global $x - z$ axes, the material 1 – 3 axes, and the wave propagation direction. For this plot $\theta_{\text{mat}} = 15^\circ$ and $\theta_{\text{wave}} = -30^\circ$. The shading shows the z -direction solid matrix velocity. This plot depicts a fast P wave, with an angular frequency of 10^4 rad/s, in the same orthotropic sandstone medium used for the inviscid convergence tests. The grid dimensions are 800×800 cells.

4.1 shows a sample plane wave solution, with the relevant axes and angles identified. The studies presented here use θ_{wave} values from 0° to 345° counterclockwise, and θ_{mat} values from 0° to 165° counterclockwise, both in steps of 15° . (When expressed in global x - z coordinates, the system matrices \mathbf{A} and \mathbf{B} simply flip sign after a rotation of $\theta_{\text{mat}} = 180^\circ$, while \mathbf{D} is unchanged, so it is not necessary to use θ_{mat} values of 180° or over.) This gives 24 different θ_{wave} values and 12 different θ_{mat} values, for a total of 288 combinations of these angles. For each $(\theta_{\text{wave}}, \theta_{\text{mat}})$ pair, plane waves are examined in each of the three families, on 100×100 , 200×200 , 400×400 , and 800×800 cell grids, for a total of 3456 different test cases. For all cases, the time step is chosen so that the CFL number is 0.9. Because the solution is smooth, no wave limiting is used for this convergence study.

For each test case, the error is measured by taking the energy norm on each cell of the difference between the numerically obtained cell value and the true solution at the cell center. The grid 1-, 2-, and max-norms are then applied to the energy norm error field to obtain an aggregate error norm. (The grid 1-norm used here is just the linear algebraic 1-norm

Table 4.1: Summary of convergence results for inviscid test cases

Incident wave	Error norm	Convergence rate			Worst R^2 value	Error on 800×800 grid	
		Best	Worst	Mean		Best	Worst
Fast P	1-norm	2.03	2.01	2.01	0.99994	2.53×10^{-5}	6.90×10^{-5}
	2-norm	2.02	2.01	2.01	0.99995	3.04×10^{-5}	8.23×10^{-5}
	Max-norm	2.02	1.96	2.00	0.99943	5.61×10^{-5}	1.76×10^{-4}
S	1-norm	2.01	2.00	2.01	1.00000	1.31×10^{-4}	3.23×10^{-4}
	2-norm	2.01	2.00	2.00	1.00000	1.49×10^{-4}	3.74×10^{-4}
	Max-norm	2.00	1.96	1.99	0.99977	2.80×10^{-4}	7.98×10^{-4}
Slow P	1-norm	1.99	1.92	1.97	0.99922	3.05×10^{-3}	1.13×10^{-2}
	2-norm	1.99	1.94	1.97	0.99953	3.44×10^{-3}	1.27×10^{-2}
	Max-norm	1.93	1.67	1.80	0.99146	8.81×10^{-3}	3.16×10^{-2}

divided by the number of grid cells; similarly, the grid 2-norm is the linear algebraic 2-norm divided by the square root of the number of grid cells.) For each combination of θ_{wave} , θ_{mat} , and wave family, a linear least-squares fit of $\log(\text{error})$ versus $\log(m)$ is performed, where m is the number of grid cells along each axis. The slope of this fit is considered to be the convergence rate of the code for this set of cases.

Table 4.1 summarizes the results of this convergence study. The convergence rates listed are the maximum, minimum, and mean over all combinations of θ_{wave} and θ_{mat} ; the worst R^2 value for the least-squares fit of $\log(\text{error})$ versus $\log(m)$ is also reported. The last two columns of Table 4.1 give the lowest and highest error for each wave in each norm on the finest grid, observed over all combinations of θ_{wave} and θ_{mat} , normalized by the energy norm of the plane wave eigenvector \mathbf{Q}_0 formed from the vectors \mathbf{V}_0 and \mathbf{T}_0 of the preceding section in order to allow a fair comparison between different cases. The solution converges at second order in all three norms for the fast P and S waves. Results are impaired for the slow P-wave because its slow propagation speed causes it to be underresolved on the coarser grids at the frequency used, but the numerical solution still attains second-order convergence in the 1-norm and 2-norm. While the error varies somewhat depending on the wave propagation and principal material directions, it does so by no more than a factor of 3-4, indicating that there are no severe grid alignment effects.

Convergence results with wave limiting here are not included here, but informal exploration suggests that for the cases above, limiting reduces the order of accuracy to about 1.9 in the 1-norm, 1.8 in the 2-norm, and 1.6 in the max-norm. This is because wave limiting tends to clip extrema in order to avoid introducing spurious oscillations. Despite the reduced order of accuracy, however, using a limiter can improve actual error in many cases, often in the 1-norm but even in the max-norm if a wave is poorly resolved or heavily




Figure 4.2: Comparison of error both with and without a limiter for the inviscid plane wave test cases with $\theta_{\text{mat}} = \theta_{\text{wave}} = 0$. The circles indicate normalized energy max-norm error, while the stars indicate normalized energy 1-norm error. The limiter is always beneficial in the 1-norm for these cases, and is also helpful in the max-norm for poorly-resolved waves.

affected by dispersive errors. Figure 4.2 shows an example of the effect of limiting, using the Monotonized Centered (MC) limiter, on the normalized energy max-norm and 1-norm errors in each of the three wave families for the inviscid test cases with $\theta_{\text{mat}} = \theta_{\text{wave}} = 0$. For each wave, the max-norm error decreases more slowly with increasing grid size when the limiter is present. The fast P wave is well resolved even on the coarsest grid, and always shows lower max-norm error without limiting. The S wave is somewhat less well-resolved, and the difference between the two curves is smaller, with the max-norm errors both with and without limiters roughly equal on the coarsest grid. Finally, the slow P wave starts out poorly resolved, and using a limiter produces lower max-norm error on all but the finest grid. The 1-norm error is lower with the limiter included for all cases. For further discussion of the benefits and drawbacks of using limiters, see LeVeque [57]. There have also been efforts to produce limiters that are compatible with higher-order methods; see for example Čada and Torrilhon [15], Liu and Tadmor [60], or the recent review by Kemm [51].




Figure 4.3: Wave speeds (meters per second) for full and reduced systems for both the materials used in this chapter with viscosity included, as a function of propagation angle. Zero propagation angle corresponds to the material principal 1-axis. The strict subcharacteristic condition (3.4) is satisfied.

4.3 Plane wave convergence results — viscous

With viscosity included, numerical solution of the equations of poroelasticity becomes substantially more challenging. The chief difficulty is that the dissipation term has its own associated time scales, independent of the computational grid. Since an appropriate time step for the wave propagation part of the system is proportional to the grid size — preferably with a CFL number near 1 — for large enough grid cell sizes the dissipation term is stiff relative to the wave propagation term. While stability is maintained by solving the dissipation term exactly, because the wave propagation and dissipation parts of the system do not commute, we can still heuristically expect problems in the operator splitting scheme if the time step is much longer than the characteristic time scale for dissipation.

Revisiting the subcharacteristic condition of Section 3.3, Figure 4.3 shows the wave speeds for the full and reduced systems as a function of propagation direction relative to the principal axes. The strict subcharacteristic condition (3.4) is satisfied for all the materials examined, although the reduced P wave speed nearly reaches the fast P wave speed for the glass/epoxy material. In fact, an even stricter condition is satisfied: the wave speeds interleave, with exactly one wave of the reduced system between each consecutive pair of waves of the full system. Based on the discussion of Section 3.3, this suggests that we will not see spurious solutions or incorrect wave speeds from the numerical solution.

To better explore the behavior of the numerical method in the presence of dissipation, a series of numerical tests were run in the same sandstone medium as the inviscid test cases, against plane wave solutions at frequencies ranging from 10 Hz to 20 kHz. (The maximum

frequency for low-frequency Biot theory to be valid in this medium is roughly 25 kHz.) For all cases, the material 1-axis was aligned with the global x -axis, and waves were set to propagate in the positive x direction. For the fast P and S waves, the domain size is chosen to be two damped wavelengths of the wave in question; since the slow P wave has a characteristic decay length (the distance over which the wave amplitude decreases by a factor of e) that is typically a fifth or less of its wavelength, the domain size for the slow P wave cases is twice the characteristic decay length instead. The domain sizes are chosen in this way so that the number of grid cells per wavelength, or per decay length, would be constant across all frequencies; this keeps the discretization error contributed from the wave propagation part of the system roughly constant for each grid size across all frequencies, helping to isolate the error caused by operator splitting. The total simulation time for the fast P and S wave cases is 1.25 cycles of the wave (a non-integer number of cycles helps avoid any possible spoofing where an unchanged solution might appear correct), while for the slow P wave cases it is 1.25 times the time for a fast P wave traveling in the x direction to cross the domain. For each combination of wave family and frequency, the solution is computed on grids of size 100×100 , 200×200 , 400×400 , and 800×800 , using both Godunov and Strang splitting.

Figure 4.4 shows the results of these tests in the same normalized energy max-norm used for the inviscid cases. There is a pronounced qualitative difference in convergence behavior depending on frequency. At low frequencies, corresponding to large grid cell sizes and long time steps, both splitting methods show first-order convergence. Starting at a step length of roughly 5-10 times the characteristic time τ_d for dissipation in the x direction, the two methods begin behaving differently, with the Godunov splitting error increasing abruptly while the Strang splitting error sweeps smoothly down to second-order convergence. This effect is most visible for the fast P wave, but can also be seen slightly in the 400×400 grid and strongly in the 800×800 grid for the S wave. Because of the choice of domain size, the slow P test cases always have a time step below the characteristic dissipation time; they display consistent first-order convergence with Godunov splitting and second-order with Strang. The qualitative shift in behavior with time step length can be understood by noting that for a time step much longer than the characteristic time, the solution operator $\exp(\mathbf{D}\Delta t)$ of the dissipative part of the system is essentially a projection operator that sets the fluid relative velocity to zero and transfers all of the fluid relative momentum into the bulk motion of the medium. The effect of this projection operator is essentially the same whether it is applied once per time step after solution of the wave propagation part of the system (Godunov splitting), or twice, both before and after wave propagation (Strang splitting).

With these results available to inform the choice of further test cases, a set of convergence studies similar to the inviscid cases in the previous subsection were conducted. Because of the qualitative difference in convergence behavior for different time step regimes, convergence studies were performed both at a point in the high-frequency convergence regime of Figure 4.4 (10 kHz), and at a low-frequency point (10 Hz). Since the slow P wave decays extremely rapidly in the presence of viscosity — typically by a factor of 10 to 100 or more per wavelength in the valid frequency range for Biot theory — the domain has a different size for the viscous slow P test cases. The high-frequency cases use square domains with side




Figure 4.4: Normalized energy max-norm error for operator splitting tests. For each splitting method (Godunov or Strang) within each subplot, the curves fall in order of increasing grid fineness, with each curve from a grid twice as fine as the curve above. The top curve is on a 100×100 grid, and the bottom is on an 800×800 . The circles indicate where the time step was less than the characteristic time scale τ_d for dissipation in the x direction; the stars indicate where the time step was greater.

Table 4.2: Summary of convergence results for viscous, high-frequency test cases

Incident wave	Error norm	Convergence rate			Worst R^2 value	Error on 800×800 grid	
		Best	Worst	Mean		Best	Worst
Fast P	1-norm	2.03	2.01	2.01	0.99996	2.61×10^{-5}	7.31×10^{-5}
	2-norm	2.03	2.01	2.01	0.99996	3.16×10^{-5}	8.79×10^{-5}
	Max-norm	2.05	2.00	2.02	0.99887	6.44×10^{-5}	1.77×10^{-4}
S	1-norm	2.01	2.01	2.01	1.00000	1.38×10^{-4}	3.39×10^{-4}
	2-norm	2.01	2.00	2.01	1.00000	1.58×10^{-4}	3.98×10^{-4}
	Max-norm	2.03	1.99	2.00	0.99978	2.99×10^{-4}	7.59×10^{-4}
Slow P	1-norm	2.01	2.00	2.01	1.00000	2.12×10^{-6}	1.74×10^{-4}
	2-norm	2.02	2.00	2.01	1.00000	2.49×10^{-6}	3.26×10^{-4}
	Max-norm	2.02	1.96	2.00	0.99988	6.53×10^{-6}	2.25×10^{-3}

length 1.2 meters for the fast P and S waves — roughly two wavelengths of the fast P wave at 10 kHz — and 5 centimeters for the slow P wave — roughly 2-5 times the characteristic decay length for this wave, depending on propagation direction in the orthotropic medium. For the low-frequency cases, the domain size is 1200 meters for the fast P and S waves, and 1 meter for the slow P wave — again, 2-4 times the characteristic decay length, which is far shorter than the wavelength. This huge disparity in domain sizes is somewhat troublesome, but it is not clear whether simulation results for a slow P wave on a domain of the size used for the fast P and S waves would be meaningful, since the solution decays over such a short distance. For practical problems, this would be an excellent opportunity for adaptive mesh refinement, to generate fine grids where and when slow P waves appear, then coarsen the grid again after they dissipate.

All of the viscous test cases run to essentially the same final times as for the frequency sweep of Figure 4.4. For the 10 kHz runs, this is $125 \mu\text{s}$ for the fast P and S waves (1.25 periods of the wave), and $10.4 \mu\text{s}$ for the slow P wave (1.25 times the time for a fast P wave to cross the domain), while for the 10 Hz runs it is 0.125 s for the fast P and S waves, and $208 \mu\text{s}$ for the slow P wave. Strang splitting is used for all cases. All other aspects of the solution, including the sets of wave propagation and material principal directions θ_{wave} and θ_{mat} as well as the method of setting the boundary conditions, are the same as for the inviscid test cases, and the solution error is measured using the same set of norms.

Table 4.2 summarizes convergence for the high-frequency viscous cases. The numerical solution again demonstrates consistent second-order convergence in all norms for the fast P and S waves at high frequency, and a similar amount of dependence of error on the wave propagation direction. For the slow P wave at high frequency, however, results are substantially different from the inviscid cases. The solution still converges at second order,

Table 4.3: Summary of convergence results for viscous, low-frequency test cases

Incident wave	Error norm	Convergence rate			Worst R^2 value	Error on 800×800 grid	
		Best	Worst	Mean		Best	Worst
Fast P	1-norm	1.57	1.10	1.26	0.99239	1.95×10^{-4}	4.24×10^{-4}
	2-norm	1.60	1.10	1.28	0.99206	2.15×10^{-4}	4.77×10^{-4}
	Max-norm	1.61	1.10	1.32	0.99014	3.90×10^{-4}	1.04×10^{-3}
S	1-norm	1.85	1.37	1.58	0.99241	4.37×10^{-4}	7.30×10^{-4}
	2-norm	1.85	1.37	1.59	0.99274	5.01×10^{-4}	8.30×10^{-4}
	Max-norm	1.83	1.36	1.58	0.98917	9.63×10^{-4}	1.92×10^{-3}
Slow P	1-norm	2.00	1.94	1.98	0.99985	1.04×10^{-6}	4.88×10^{-5}
	2-norm	2.03	1.93	1.99	0.99982	1.08×10^{-6}	7.40×10^{-5}
	Max-norm	2.09	1.84	1.98	0.99795	2.40×10^{-6}	4.59×10^{-4}

since the solution is well-resolved on the grids used here, but there is now a factor of several hundred difference between the maximum and minimum error at a single grid size. Close examination of the error for individual cases shows that it is primarily a function of the offset $\theta_{\text{wave}} - \theta_{\text{mat}}$; for a fixed value of $\theta_{\text{wave}} - \theta_{\text{mat}}$ and a fixed grid size, the error is similar across all values of θ_{mat} . This indicates that the large variation in error is a effect of the alignment of the wavefront relative to the principal material axes, rather than a grid alignment effect. The likely cause is the substantial difference in the characteristic decay times between the 1 and 3 axes of the material — the decay time in the 1 direction is 5.95 microseconds, while in the 3 direction it is 1.82 microseconds. This large variation in decay time causes a large variation in the operator splitting error. In addition, the characteristic decay length is substantially shorter in the 3 direction, causing the solution magnitude to be larger at the “upstream” (opposite the propagation direction) edge of the domain relative to the value at the center against which the error is normalized; the larger solution magnitude naturally results in a larger error.

Table 4.3 shows the results for the low-frequency viscous test cases. The fast P and S waves again show only a weak dependence of error on wave propagation and principal material direction, but their convergence rates are substantially degraded, just as in the low-frequency range of Figure 4.4. Convergence of the fast P wave is roughly first-order in the worst case, due to the long time step relative to the characteristic decay times. Surprisingly, results for the S wave are only slightly better in the worst case, likely due to the much shorter characteristic decay time in the 3 direction. If the grid were further refined, the numerical solution would presumably reach a second-order convergence regime as the time step approached the characteristic decay time, but to have a time step similar to the shortest decay time at a CFL number of 0.9 would require a grid cell size of roughly

13 cm — resulting in a $90,000 \times 90,000$ cell grid on the 1200 m square domain! The low-frequency slow P wave cases, by contrast, show the same strong dependence of error on the alignment of the wave direction with the principal material direction as the high-frequency cases, but because the domain size is much smaller and the time step much shorter, these cases show consistent second-order convergence.

As a final comment on convergence, while the rate of reduction of error with decreasing mesh size is poorer for the low-frequency fast P and S wave cases, this may not be a problem in practice unless very high accuracy is desired. For both frequency ranges investigated, the relative error in all three norms on the 800×800 grid never exceeds 1.04×10^{-3} for the fast P wave, or 1.92×10^{-3} for the S wave — despite the numerical difficulties encountered, the waves are still well-resolved.

4.4 *Single-material point source results*

With the accuracy of the method characterized for simple plane wave solutions, the discussion may move on to more interesting problems. The CLAWPACK solution will first be compared against the results of de la Puente et al. [29] and Carcione [17] for a point source in a uniform orthotropic medium. These simulations use the orthotropic sandstone and glass/epoxy media described in Table A.1; in both cases, the material 1 axis coincides with the x axis. All test cases start with initial condition $\mathbf{Q}(x, z, 0) = 0$, and the domain is excited by a point source with a Ricker wavelet profile having peak frequency $f_{\text{src}} = 3730$ Hz for sandstone and 3135 Hz for glass/epoxy, acting with peak intensity $+1 \text{ Pa} \cdot \text{m}^2/\text{s}$ on the vertical normal stress τ_{zz} and $-1 \text{ Pa} \cdot \text{m}^2/\text{s}$ on fluid pressure. The peak of the wavelet occurs 0.4 ms after the start of the simulation. For each case, a square domain is used with dimensions 18.7 m on a side; the point source is placed at the center of the domain. The simulation time span is 1.56 ms for the sandstone medium, and 1.80 ms for glass/epoxy. Results are calculated both with and without viscosity present.

All of these simulations are carried out on a uniform 501×501 cell grid. The odd number of grid cells in each axis allows the point source to be applied to a single grid cell. A grid resolution this high is necessary to resolve the slow P wave well; as in the plane wave test cases, the fast P and S waves are well-resolved on substantially coarser grids. The point source is implemented numerically as part of the source step in the operator splitting scheme; since the point source acts on stress variables, and the viscous dissipation acts on velocity variables, it does not matter which is applied first. The CFL number for all simulations is 0.9, resulting in 279 time steps being taken for the sandstone case and 282 for the glass/epoxy. Limiters are enabled for these solutions — all waves are limited using the monotonized centered (MC) limiter, since even though most of the solution is smooth, without a limiter the lack of smoothness at the source point produces substantial spurious oscillations in the solution. Figure 4.5 shows the results of these simulations. These figures correspond to Figures 6 and 7 of de la Puente et al. [29], or Figures 5 and 7 of Carcione [17], and are in agreement with them.




Figure 4.5: Results for point source test cases. Top: sandstone; bottom: glass/epoxy; left to right: v_x without viscosity, v_x with viscosity, v_z without viscosity, v_z with viscosity.

4.5 Heterogeneous domain results

The results for purely proelastic problems on rectilinear grids conclude with a wave reflection and interconversion problem that demonstrates the ability of the method to model material interfaces, as well as the benefits of using adaptive mesh refinement — a feature of the AMRCLAW variant of CLAWPACK that has not been used thus far.

The final test case is a large-scale inviscid problem with a bed of isotropic shale overlying isotropic sandstone, with material properties given in Table A.1. The problem domain is the rectangle $[0, 1500 \text{ m}] \times [0, 1400 \text{ m}]$ in the x - z plane, with the boundary between the materials at $z = 700 \text{ m}$. There is a point source at $(x, z) = (750 \text{ m}, 900 \text{ m})$, again with a Ricker wavelet profile in time, with peak frequency 50 Hz. The source acts on the z -direction normal stress and fluid pressure with peak intensities $+2.3 \times 10^{13} \text{ Pa} \cdot \text{m}^2/\text{s}$ and $-2.3 \times 10^{13} \text{ Pa} \cdot \text{m}^2/\text{s}$, respectively, similarly to the previous test cases in homogeneous domains. This source magnitude is chosen to roughly match the magnitude of the response shown by de la Puente et al. [29]. The peak of the source is delayed 40 ms after the start of the simulation, and the total duration of the run is 0.5 s. In addition to time-snapshots of the solution at particular instants, solution time histories are also recorded at three “gauges,” located at $(x_1, z_1) = (950 \text{ m}, 750 \text{ m})$, $(x_2, z_2) = (950 \text{ m}, 650 \text{ m})$, and $(x_3, z_3) = (950 \text{ m}, 500 \text{ m})$.

For the adaptively refined simulation, the coarsest-level computational grid was 75×70 cells in size, giving square cells 20 m on a side. Two additional levels of refinement are used

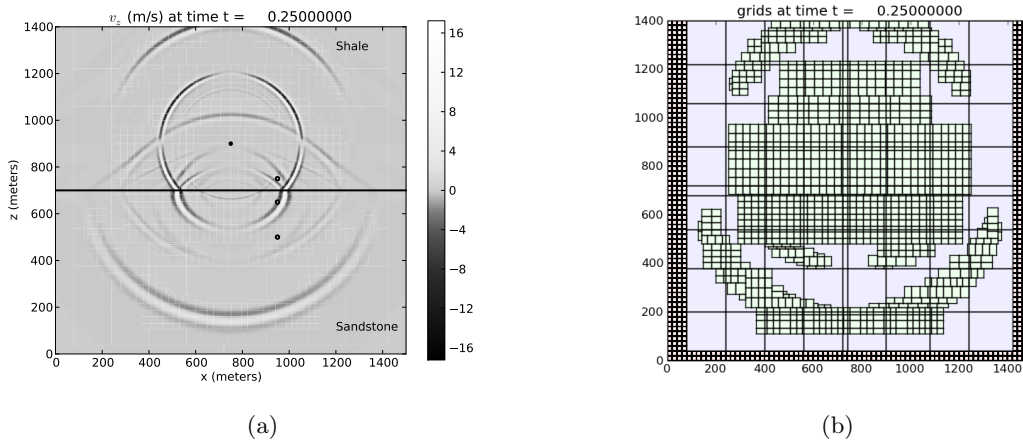


Figure 4.6: Left: snapshot of z -direction solid velocity 0.25 seconds after the start of the simulation. The source location is marked with a solid black dot, while the gauges are marked with white-centered black dots. The gauges are numbered from top to bottom. Right: AMR grids at this time point. Individual cells are drawn on the coarsest AMR level, but only grid outlines are shown on finer levels.

on this grid, the first at a factor of 4, and the second at further factor of 6, so that cells on the finest grids were 0.83 m on a side. The code flagged a cell for mesh refinement when the energy norm (using the material properties of that cell) of the difference $\Delta \mathbf{Q}$ between its state vector and that of any adjacent cell exceeded $32.5 \text{ J}^{1/2}/\text{m}^{3/2}$, with the exception of the rectangle $[700 \text{ m}, 1000 \text{ m}] \times [450 \text{ m}, 950 \text{ m}]$, where the threshold for refinement is lowered to $3.25 \text{ J}^{1/2}/\text{m}^{3/2}$ in order to improve accuracy at the gauges. These tolerances are chosen empirically based on the observed magnitude of the waves in the simulation. This refinement criterion is a generalization of the typical AMRCLAW approach of refining based on the difference between the solution values in neighboring cells, which has been used successfully on many problems; AMRCLAW makes it easy to set alternate user-specified refinement criteria if desired, and also offers automatic error estimation via Richardson extrapolation. Besides refinement based on the solution field, the source location is also flagged for refinement to the finest level available whenever the source intensity is greater than about 10^{-9} of peak. Since the grid size in each direction is an even number on all but the coarsest grids, the source is distributed over the four cells closest to its location using a bilinear weighting. The MC limiter is again used for all waves in this problem.

Figure 4.6 shows a snapshot of the z -direction solid velocity 0.25 s after the start of the simulation, analogous to Figure 9(a) of de la Puente et al. [29], along with the AMR grids at this time. The solid black dot indicates the source location, and the white-centered black dots indicate the gauge locations. Because the eigenvectors associated with each wave family are different in the two materials, when a wave impinges on the material interface it

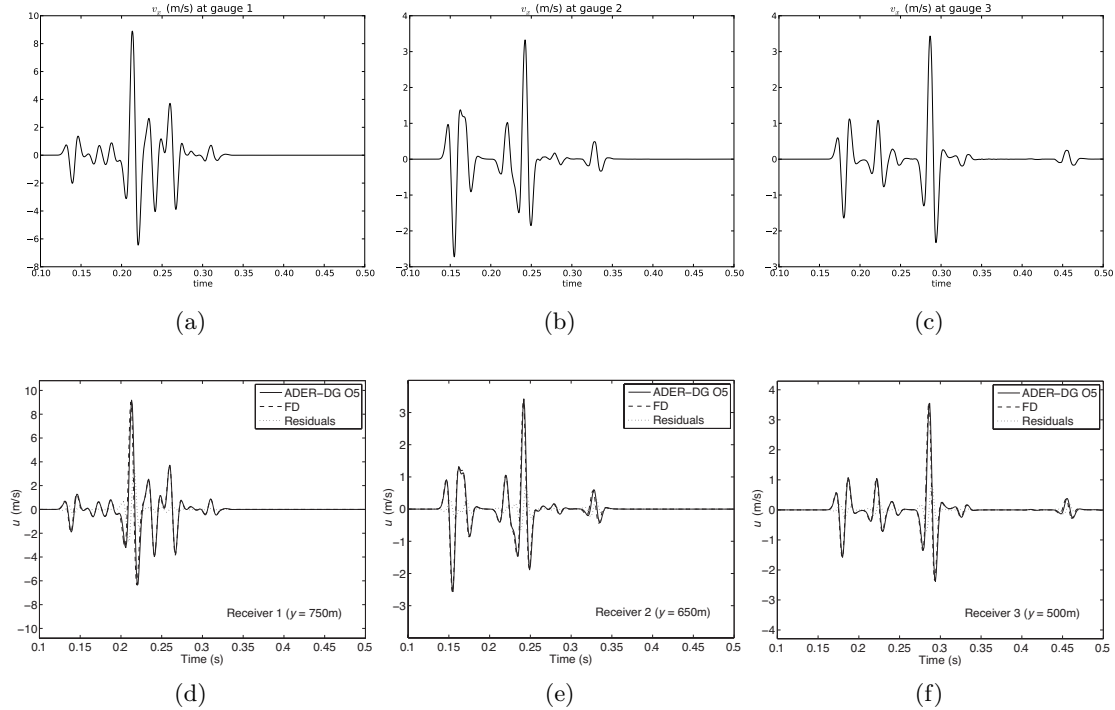


Figure 4.7: Time-histories of x -direction solid velocity at the gauges indicated in Figure 4.6a. Top: CLAWPACK; bottom: de la Puente et al. [29]. Left to right: gauge 1, at $(x, z) = (950 \text{ m}, 750 \text{ m})$; gauge 2, at $(950 \text{ m}, 650 \text{ m})$; gauge 3, at $(950 \text{ m}, 500 \text{ m})$. Results of de la Puente et al. are reproduced in accordance with the policies of the publishing journal; the ADER-DG and FD curves in these plots are discontinuous Galerkin and finite difference results, respectively.

produces reflected and transmitted waves in each of the three families; this results in a rich and complex solution structure. In addition to the reflected and transmitted waves, head waves also appear, where a wave in the lower half of the domain excites a wave in a slower family in the upper half. This results in a straight wavefront, rather than a curved one. At this point in the simulation the level 2 AMR grids have expanded to cover most of the domain, but the level 3 grids are concentrated around the wavefronts. Figures 4.7 and 4.8 show the time histories of the solid x and z velocities at the three gauges. The results are in generally good agreement with Figure 10 of de la Puente et al. [29], also shown in Figures 4.7 and 4.8, although the peaks of slow P wave event at $t = 0.45 \text{ s}$ at gauge 3 are clipped in the CLAWPACK simulation because of the limiter, and the magnitude of the second large excursion in vertical velocity at gauge 3 in the solution here seems somewhat less.

While the variety of different wave speeds and reflected/transmitted waves in this problem make adaptive mesh refinement less useful than is typically the case — by halfway

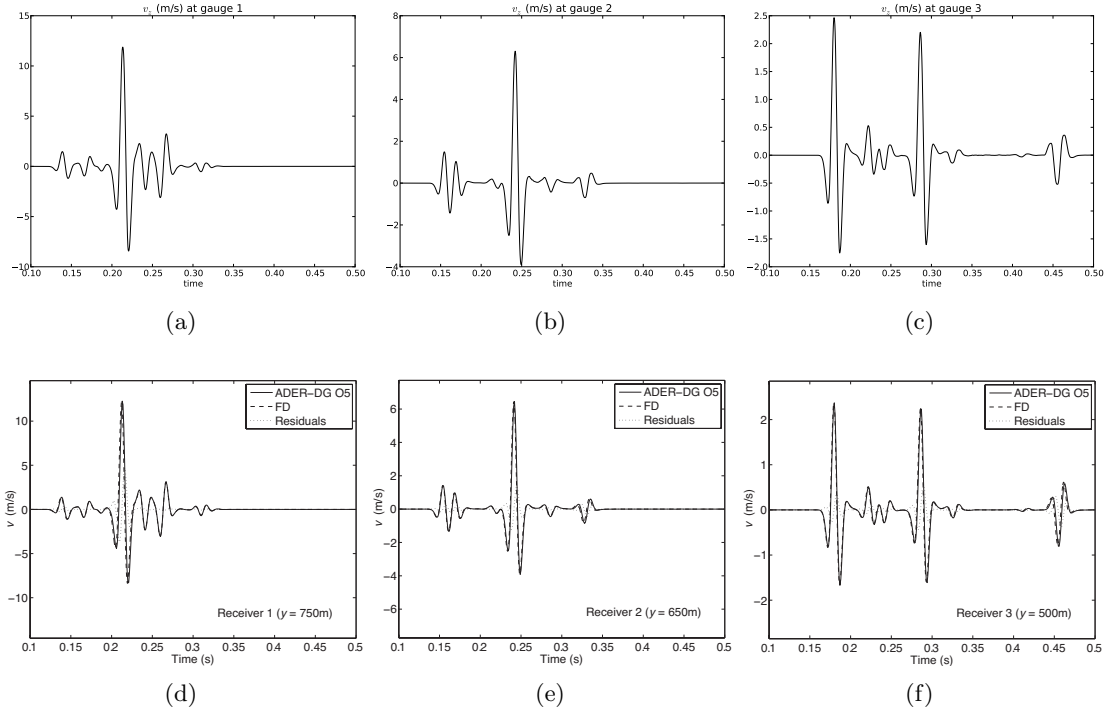


Figure 4.8: Time-histories of z -direction solid velocity at the gauges indicated in Figure 4.6a. Top: CLAWPACK; bottom: de la Puente et al. [29]. Left to right: gauge 1, at $(x, z) = (950 \text{ m}, 750 \text{ m})$; gauge 2, at $(950 \text{ m}, 650 \text{ m})$; gauge 3, at $(950 \text{ m}, 500 \text{ m})$. Results of de la Puente et al. are reproduced in accordance with the policies of the publishing journal; the ADER-DG and FD curves in these plots are discontinuous Galerkin and finite difference results, respectively.

through the simulation time, Figure 4.6b shows a large portion of the domain refined at the finest level, because there are wavefronts present throughout the domain — AMR still results in a substantial savings in computation time. On an Amazon EC2 Cluster 8XL instance, running with 32 OpenMP threads, these results took 20 minutes 31 seconds to obtain, whereas a uniformly refined grid with the same cell size as the finest AMR grids took 47 minutes 19 seconds and produced no significant change in the solution. (Both times are the average of two runs; each pair differed by 4 seconds or less.) The large number of hardware threads available on this type of EC2 instance is the reason why there are so many separate fine grids in Figure 4.6b — AMRCLAW uses a coarse-grained parallelization strategy, with each grid at each time step processed by a single thread, which means that many grids must be present in order to take full advantage of highly parallel computers. The number of grids used at each refinement level is indirectly controlled by setting the maximum size of the individual grids; for the AMR computation shown above, grids were

allowed to extend no more than 60 cells in any direction. Having a very large number of small grids (1047 level 3 grids in the figure) also eases load balancing between threads, since most grids are of similar size, and each thread processes many grids.

Chapter 5

METHODS FOR POROELASTIC-FLUID SYSTEMS ON TWO-DIMENSIONAL MAPPED GRIDS

In order to be able to model geometries that do not lend themselves to the straight lines and 90 degree angles of a rectilinear grid, let us now turn to logically rectangular mapped grids. Compared to unstructured grids, mapped grids have the advantage of simpler data structures and less computational overhead, although creating the desired mapping function is not always trivial. Mapped grids also combine well with finite volume methods, since the finite volume solution tends to maintain good quality even in the face of a mapping that severely distorts the grid. However, when modeling anisotropic poroelastic media, having to deal with cell interfaces that may be oriented in any arbitrary direction means that the simple compute-and-cache method used to solve the Riemann problem in chapters 3 and 4 no longer suffices. Section 5.2 discusses a Riemann solution process designed to function efficiently in this context, and to also incorporate interface conditions such as (2.103) and (2.105).

The introduction of fluid media into the numerical simulation causes an additional difficulty that must be addressed. At an interface between a poroelastic medium and a fluid, the classical formulation of a high-resolution finite volume method can be qualitatively incorrect — the classical transverse and second-order correction fluxes result in poroelastic variables such as skeleton stress and solid velocity being carried into the fluid, where they make no sense. This is discussed further in Section 5.3, and the method is reformulated to produce qualitatively correct solutions, though the best way to fully adapt high-resolution finite volume methods, including generalizations of the classical limiters and second-order correction fluxes, is still an open question.

The remaining part of this chapter briefly discusses the implementation of the source term in the poroelasticity system and the CLAWPACK software framework in which the numerical solution is implemented.

5.1 Mapped grids

For a mapped grid approach, start with a uniform rectangular grid in the computational coordinates, denoted ξ_1 and ξ_2 . Then apply a mapping function $\mathbf{X}(\xi_1, \xi_2)$ to obtain the grid in physical coordinates (x, z) . This mapping function is typically chosen so that the grid boundaries or interior grid lines follow some feature of interest in the problem, although once the grid mapping is formulated, the actual cells are taken to be quadrilaterals with straight sides, whose vertices are obtained using the mapping function. Each grid cell ij has an associated capacity κ_{ij} , which is the ratio of the area of the cell in physical coordinates to its area $\Delta\xi_1 \Delta\xi_2$ in computational coordinates, and each cell interface has an associated unit normal vector \mathbf{n} pointing in the positive grid direction. Figure 5.1 shows an example

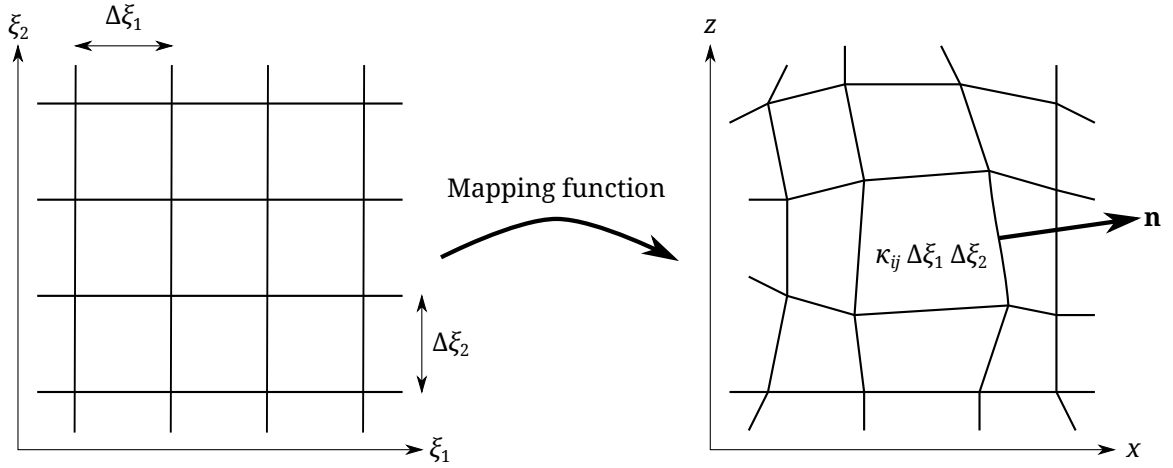


Figure 5.1: Example of a mapped grid, showing computational space at the left and physical space at the right.

of a mapped grid. Finite volume methods on mapped grids are discussed in greater detail in Chapter 23 of [57].

5.2 Riemann problems on mapped grids

One of the most critical parts of any high-resolution finite volume code is the solution of the Riemann problem. Solving the Riemann problem for poroelasticity without the caching approach of Chapter 3 at first appears challenging due to the size and complexity of the hyperbolic part of the system (2.43). However, by taking advantage of the structure of the system, a straightforward and fairly efficient solution is possible.

5.2.1 Eigenvalues and eigenvectors

The first order of business in solving the Riemann problem is finding the eigenvectors of $\check{\mathbf{A}} = n_x \mathbf{A} + n_z \mathbf{B}$ corresponding to the propagating waves, and the corresponding eigenvalues giving the wave speeds. For linear acoustics, this eigensystem is simple: the eigenvalues are $\pm \sqrt{K_f / \rho_f}$, and the eigenvectors may be easily verified as

$$\begin{aligned} \mathbf{r}_{\text{acoustic, left}} &= (-Z_f \ 0 \ 0 \ 0 \ 0 \ 0 \ n_x \ n_z)^T \quad (\text{left-going wave}) \\ \mathbf{r}_{\text{acoustic, right}} &= (Z_f \ 0 \ 0 \ 0 \ 0 \ 0 \ n_x \ n_z)^T \quad (\text{right-going wave}). \end{aligned} \quad (5.1)$$

For poroelasticity the eigensystem is substantially more complicated, but it is possible to exploit the block structure, symmetry, and simple null space of the poroelastic system to reduce the eigenproblem from an 8×8 non-symmetric to a 3×3 real symmetric one. Let (n_1, n_3) be the components of the unit normal to the Riemann solve interface, measured in

the material principal coordinates. Then, writing the first-order system (2.43) in principal coordinates, $\check{\mathbf{A}}$ is

$$\check{\mathbf{A}} = n_1 \mathbf{A} + n_3 \mathbf{B}, \quad (5.2)$$

where \mathbf{A} and \mathbf{B} are exactly as given in (2.45), (2.47), and (2.48). Furthermore, from the properties of \mathbf{A} and \mathbf{B} , we know that $\mathbf{E}\check{\mathbf{A}}$ is symmetric; defining $\check{\mathbf{A}}_{sv} := n_1 \mathbf{A}_{sv} + n_3 \mathbf{B}_{sv}$ and $\check{\mathbf{A}}_{vs} := n_1 \mathbf{A}_{vs} + n_3 \mathbf{B}_{vs}$, we get $\mathbf{E}_s \check{\mathbf{A}}_{sv} = (\mathbf{E}_v \check{\mathbf{A}}_{vs})^T$.

In order to see how this is useful, let λ and \mathbf{r} be an eigenpair of $\check{\mathbf{A}}$, so that $\check{\mathbf{A}}\mathbf{r} = \lambda\mathbf{r}$. Since \mathbf{E} is nonsingular, this is equivalent to the generalized eigenproblem

$$\mathbf{E}\check{\mathbf{A}}\mathbf{r} = \lambda\mathbf{E}\mathbf{r}. \quad (5.3)$$

Dividing the eigenvector into stress and velocity blocks, $\mathbf{r} =: (\mathbf{r}_s^T \quad \mathbf{r}_v^T)^T$, and referring back to the generalized eigenproblem formulation of 2.1.9, the second equation of (2.65) can be rewritten as $(\mathbf{E}_s \check{\mathbf{A}}_{sv})^T \mathbf{r}_s = \check{\mathbf{A}}_{sv}^T \mathbf{E}_s \mathbf{r}_s = \lambda \mathbf{E}_v \mathbf{r}_v$. Multiplying the first equation on the left by $\check{\mathbf{A}}_{sv}^T$ gives a 4×4 symmetric-definite eigenproblem for λ^2 and \mathbf{r}_v :

$$\check{\mathbf{A}}_{sv}^T \mathbf{E}_s \check{\mathbf{A}}_{sv} \mathbf{r}_v = \lambda^2 \mathbf{E}_v \mathbf{r}_v. \quad (5.4)$$

Aside from reducing the dimension of the eigenproblem, this also illustrates that the eigenvalues of this system come in positive and negative pairs — if λ is an eigenvalue, then so is $-\lambda$.

While the reduction from an 8×8 asymmetric ordinary eigenproblem to a 4×4 symmetric-definite generalized eigenproblem is already a substantial gain, further improvement is possible. First, factorize \mathbf{E}_v as $\mathbf{E}_v = \mathbf{L}\mathbf{L}^T$, where \mathbf{L} is upper-triangular; such an \mathbf{L} can be found straightforwardly as

$$\mathbf{L} = \begin{pmatrix} \sqrt{\frac{\Delta_1}{m_1}} & 0 & \frac{\rho_f}{\sqrt{m_1}} & 0 \\ 0 & \sqrt{\frac{\Delta_3}{m_3}} & 0 & \frac{\rho_f}{\sqrt{m_3}} \\ 0 & 0 & \sqrt{m_1} & 0 \\ 0 & 0 & 0 & \sqrt{m_3} \end{pmatrix}. \quad (5.5)$$

(Many factorizations of \mathbf{E}_v are possible, but the sparse upper-triangular structure of this \mathbf{L} will be useful later.) Making the variable substitution $\mathbf{r}_v = \mathbf{L}^{-T}\mathbf{y}$, then multiplying from the left by \mathbf{L}^{-1} , gives the 4×4 symmetric eigenproblem

$$\mathbf{M}\mathbf{y} = \lambda^2\mathbf{y}, \quad \mathbf{M} := \mathbf{L}^{-1} \check{\mathbf{A}}_{sv}^T \mathbf{E}_s \check{\mathbf{A}}_{sv} \mathbf{L}^{-T}. \quad (5.6)$$

The way the matrix \mathbf{M} is defined offers an additional opportunity to reduce the amount of computation performed during the Riemann solve. Substituting $\check{\mathbf{A}}_{sv} = n_1 \mathbf{A}_{sv} + n_3 \mathbf{B}_{sv}$ into the definition of \mathbf{M} , the matrix can be written as a quadratic form in n_1 and n_3 ,

$$\mathbf{M} = n_1^2 \mathbf{M}_{11} + n_1 n_3 \mathbf{M}_{13} + n_3^2 \mathbf{M}_{33}, \quad (5.7)$$

where

$$\begin{aligned} \mathbf{M}_{11} &= \mathbf{L}^{-1} \mathbf{A}_{sv}^T \mathbf{E}_s \mathbf{A}_{sv} \mathbf{L}^{-T} \\ \mathbf{M}_{13} &= \mathbf{L}^{-1} \mathbf{A}_{sv}^T \mathbf{E}_s \mathbf{B}_{sv} \mathbf{L}^{-T} + \mathbf{L}^{-1} \mathbf{B}_{sv}^T \mathbf{E}_s \mathbf{A}_{sv} \mathbf{L}^{-T} \\ \mathbf{M}_{33} &= \mathbf{L}^{-1} \mathbf{B}_{sv}^T \mathbf{E}_s \mathbf{B}_{sv} \mathbf{L}^{-T}. \end{aligned} \quad (5.8)$$

Since the matrices \mathbf{M}_{11} , \mathbf{M}_{13} , and \mathbf{M}_{33} do not depend on the interface direction but only on the properties of the medium, they can be computed once for each material and retrieved when needed to form the full \mathbf{M} .

Finally, the dimension of the eigenproblem can be reduced even further by finding one of the eigenvectors explicitly and removing it from the computation. It is easy to verify that $\mathbf{r}_{v0} = (0 \ 0 \ -n_3 \ n_1)^T$ satisfies $\check{\mathbf{A}}_{sv}\mathbf{r}_{v0} = 0$; physically, \mathbf{r}_{v0} corresponds to fluid flow parallel to the Riemann problem interface, which produces no propagating waves in inviscid Biot theory. This is in fact a basis for the entire null space of $\check{\mathbf{A}}_{sv}$, since the other three eigenvectors correspond to the three propagating waves with nonzero wave speeds λ . Thus $\mathbf{y}_0 := \mathbf{L}^T\mathbf{r}_{v0}/\|\mathbf{L}^T\mathbf{r}_{v0}\|_2$ is a unit null vector of \mathbf{M} . Now construct a 4×3 matrix \mathbf{Z} whose columns are orthonormal and all orthogonal to \mathbf{y}_0 . (This is where the sparse upper-triangular structure of \mathbf{L} comes in handy, since it means that the first two components of \mathbf{y}_0 are zero, which makes this matrix easy to construct.) The matrix $\mathbf{Y} := (\mathbf{Z} \ \mathbf{y}_0)$ is then an orthonormal matrix. Making the new variable substitution $\mathbf{y} = \mathbf{Y}(\mathbf{u}^T \ u_0)^T$ in (5.6) and multiplying from the left by \mathbf{Y}^T , we then get

$$\mathbf{Y}^T\mathbf{M}\mathbf{Y} \begin{pmatrix} \mathbf{u} \\ u_0 \end{pmatrix} = \begin{pmatrix} \mathbf{Z}^T\mathbf{M}\mathbf{Z} & 0 \\ 0 & 0 \end{pmatrix} \begin{pmatrix} \mathbf{u} \\ u_0 \end{pmatrix} = \lambda^2 \begin{pmatrix} \mathbf{u} \\ u_0 \end{pmatrix}. \quad (5.9)$$

Since we are primarily concerned with propagating waves in the Riemann solution, we can ignore the possibility of $\lambda = 0$, and need only find eigenvalues and eigenvectors of the 3×3 symmetric matrix $\mathbf{M}' := \mathbf{Z}^T\mathbf{M}\mathbf{Z}$. The classic QR algorithm with Wilkinson shifts is used for this, following the implementation outlined by Golub and van Loan [42], which has the advantages of being fairly straightforward to implement and computing all the eigenvectors at once.

The QR algorithm also has another beneficial property for the Riemann solution process. It naturally produces eigenvectors with unit 2-norm, which is useful in the orthonormalization of the eigenvectors. Backing out the variable transformations above gives $\mathbf{r}_v = \mathbf{L}^{-T}\mathbf{Z}\mathbf{u}$, so $\mathbf{r}_v^T\mathbf{E}_v\mathbf{r}_v = \mathbf{u}^T\mathbf{u}$. Dividing each eigenvector \mathbf{u} by $\sqrt{2}$ after computing it with the QR algorithm gives $\mathbf{r}_v^T\mathbf{E}_v\mathbf{r}_v = \frac{1}{2}$, and by the energy equipartition result (2.68), $\mathbf{r}^T\mathbf{E}\mathbf{r} = \mathbf{r}_s^T\mathbf{E}_s\mathbf{r}_s + \mathbf{r}_v^T\mathbf{E}_v\mathbf{r}_v = 2\mathbf{r}_v^T\mathbf{E}_v\mathbf{r}_v = 1$. This energy orthonormalization is helpful for calculating wave strengths in the common special case of a Riemann problem between identical poroelastic materials.

Once λ^2 and \mathbf{r}_v have been computed, λ is obtained as either the positive or negative square root of λ^2 ; negative λ values correspond to left-going waves, while positive λ values indicate right-going waves. Since the null space of $\check{\mathbf{A}}_{sv}$ is excluded when calculating these eigensolutions, no λ thus obtained will be zero. Finally, (2.64) yields $\mathbf{r}_s = \check{\mathbf{A}}_{sv}\mathbf{r}_v/\lambda$, giving all the components of the eigenvector \mathbf{r} measured in the principal axes of the material; the vector \mathbf{r} is then transformed into global computational axes to solve the Riemann problem.

As a note on implementation, the Riemann solver code caches the eigensystems computed for each Riemann problem, and checks whether the next Riemann problem has the same input data to the eigensolves; if so, the previous eigensystem is re-used. This speeds up execution in the common special case of large sections of the domain that have straight grid lines and uniform composition.

5.2.2 Solution between identical materials

The solution of the Riemann problem for linear acoustics with identical media on either side of the interface is simple and well-known, and will not be repeated here. The interested reader may refer to [57], as one of many sources.

Solution of the Riemann problem for poroelasticity with identical materials on either side of the interface is more complex, but it can be simplified greatly by taking advantage of the structure of the system. In particular, as shown in Section 2.1.7, the eigenvectors of the poroelasticity system are orthogonal with respect to the energy matrix \mathbf{E} , and in fact the eigensolution process detailed in the previous section produces eigenvectors that are orthonormal with respect to \mathbf{E} . Since no special interface condition is applied at cell interfaces within a homogeneous medium, to solve the Riemann problem it is only necessary to find the vector of wave strengths $\boldsymbol{\beta}$ such that

$$\mathbf{R}\boldsymbol{\beta} = \mathbf{Q}_r - \mathbf{Q}_l, \quad (5.10)$$

where \mathbf{R} is the matrix of eigenvectors of $\check{\mathbf{A}}$. By the orthonormality of the eigenvectors with respect to \mathbf{E} , multiplying from the left by $\mathbf{R}^T\mathbf{E}$ gives the simple result

$$\boldsymbol{\beta} = \mathbf{R}^T\mathbf{E}(\mathbf{Q}_r - \mathbf{Q}_l). \quad (5.11)$$

5.2.3 Solution between different materials

At an interface between different materials, the Riemann solution process becomes more complex. For a pair of fluids with different properties, the solution process of the acoustic Riemann problem is still well-known, and is covered for instance in [57]. For a pair of distinct poroelastic media, or a poroelastic medium and a fluid, an interface condition must be satisfied, either (2.103) or (2.105).

To solve the Riemann problem in this context, first note that both interface conditions can be cast in the form

$$\mathbf{C}_l \lim_{x' \rightarrow 0^-} \mathbf{Q} = \mathbf{C}_r \lim_{x' \rightarrow 0^+} \mathbf{Q} \quad (5.12)$$

for some matrices \mathbf{C}_l and \mathbf{C}_r . Here x' is the signed normal distance from the interface, and is positive on the right side. The solution within the left and right media is still a set of discontinuities in \mathbf{Q} proportional to the eigenvectors of the $\check{\mathbf{A}}$ matrices in the respective media, propagating at speeds equal to the corresponding eigenvalues, so $\lim_{x' \rightarrow 0^-} \mathbf{Q}$ and $\lim_{x' \rightarrow 0^+} \mathbf{Q}$ can be related to the left and right states \mathbf{Q}_l and \mathbf{Q}_r by

$$\lim_{x' \rightarrow 0^-} \mathbf{Q} = \mathbf{Q}_l + \sum_{\text{left}} \beta_{li} \mathbf{r}_{li}, \quad \lim_{x' \rightarrow 0^+} \mathbf{Q} = \mathbf{Q}_r - \sum_{\text{right}} \beta_{ri} \mathbf{r}_{ri}. \quad (5.13)$$

Here β_{li} and β_{ri} are the strengths of the left-going and right-going waves, \mathbf{r}_{li} and \mathbf{r}_{ri} are the corresponding eigenvectors, and the sums are over only the left-going and right-going waves, respectively. Substituting this into (5.12) and rearranging gives a linear system for the wave strengths,

$$(\mathbf{C}_l \mathbf{R}_l \quad \mathbf{C}_r \mathbf{R}_r) \begin{pmatrix} \boldsymbol{\beta}_l \\ \boldsymbol{\beta}_r \end{pmatrix} = \mathbf{C}_r \mathbf{Q}_r - \mathbf{C}_l \mathbf{Q}_l, \quad (5.14)$$

where \mathbf{R}_l and \mathbf{R}_r are the matrices of left-going and right-going wave eigenvectors, respectively. Note that while it is possible to formulate this linear system, there is no *a priori* guarantee that it has a solution; at a minimum, there must be exactly as many equations in the interface condition (5.12) as there are propagating waves. Solution of this system has, however, not been a problem for any of the cases considered here, and heuristically it may be expected to have a solution whenever the Riemann problem is well-posed, since if any solution of the Riemann problem at the interface exists, obtaining it is equivalent to solving (5.14).

Now all that remains is to explicitly write the matrices \mathbf{C}_l and \mathbf{C}_r corresponding to (2.103) and (2.105). Condition (2.105) will be addressed first; taking the left medium to be poroelastic, a straightforward component-by-component accounting gives

$$\mathbf{C}_{l,\text{poro-fluid}} = \begin{pmatrix} 0 & 0 & 0 & 0 & n_x & n_z & n_x & n_z \\ 0 & n_x & 0 & n_z & 0 & 0 & 0 & 0 \\ 0 & 0 & n_z & n_x & 0 & 0 & 0 & 0 \\ \eta_d & 0 & 0 & 0 & 0 & 0 & -Z_f(1-\eta_d)n_x & -Z_f(1-\eta_d)n_z \end{pmatrix} \quad (5.15)$$

$$\mathbf{C}_{r,\text{poro-fluid}} = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & n_x & n_z \\ -n_x & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ -n_z & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \eta_d & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}. \quad (5.16)$$

If the poroelastic material is on the right side of the interface, the appropriate coefficient matrices can be obtained by simply exchanging the subscripts l and r , and negating \mathbf{n} .

Writing the matrices \mathbf{C}_l and \mathbf{C}_r for (2.103) is less straightforward, due to the ambiguity in $\widehat{\mathbf{q}} \cdot \widehat{\mathbf{n}}$. This quantity is equal to both $\mathbf{q}_l \cdot \mathbf{n}$ and $\mathbf{q}_r \cdot \mathbf{n}$; in exact arithmetic it is irrelevant which is chosen, but it is not obvious how best to define $\widehat{\mathbf{q}} \cdot \widehat{\mathbf{n}}$ for numerical solution. Faced with this ambiguity, let us set $\widehat{\mathbf{q}} \cdot \widehat{\mathbf{n}} = (1 - \zeta)\mathbf{q}_l \cdot \mathbf{n} + \zeta\mathbf{q}_r \cdot \mathbf{n}$, where ζ is a free parameter used to improve the conditioning of (5.14). With this choice of $\widehat{\mathbf{q}} \cdot \widehat{\mathbf{n}}$, \mathbf{C}_l and \mathbf{C}_r become

$$\mathbf{C}_{l,\text{poro-poro}} = \begin{pmatrix} 0 & n_x & 0 & n_z & 0 & 0 & 0 & 0 \\ 0 & 0 & n_z & n_x & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & n_x & n_z \\ \eta_d & 0 & 0 & 0 & 0 & 0 & -(1-\zeta)Z_f(1-\eta_d)n_x & -(1-\zeta)Z_f(1-\eta_d)n_z \end{pmatrix} \quad (5.17)$$

$$\mathbf{C}_{r,\text{poro-poro}} = \begin{pmatrix} 0 & n_x & 0 & n_z & 0 & 0 & 0 & 0 \\ 0 & 0 & n_z & n_x & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & n_x & n_z \\ \eta_d & 0 & 0 & 0 & 0 & 0 & \zeta Z_f(1-\eta_d)n_x & \zeta Z_f(1-\eta_d)n_z \end{pmatrix}. \quad (5.18)$$

Figure 5.2 shows the variation of the condition number of the coefficient matrix in (5.14) as a function of ζ for various η_d values and various pairs of materials. The choice of ζ is

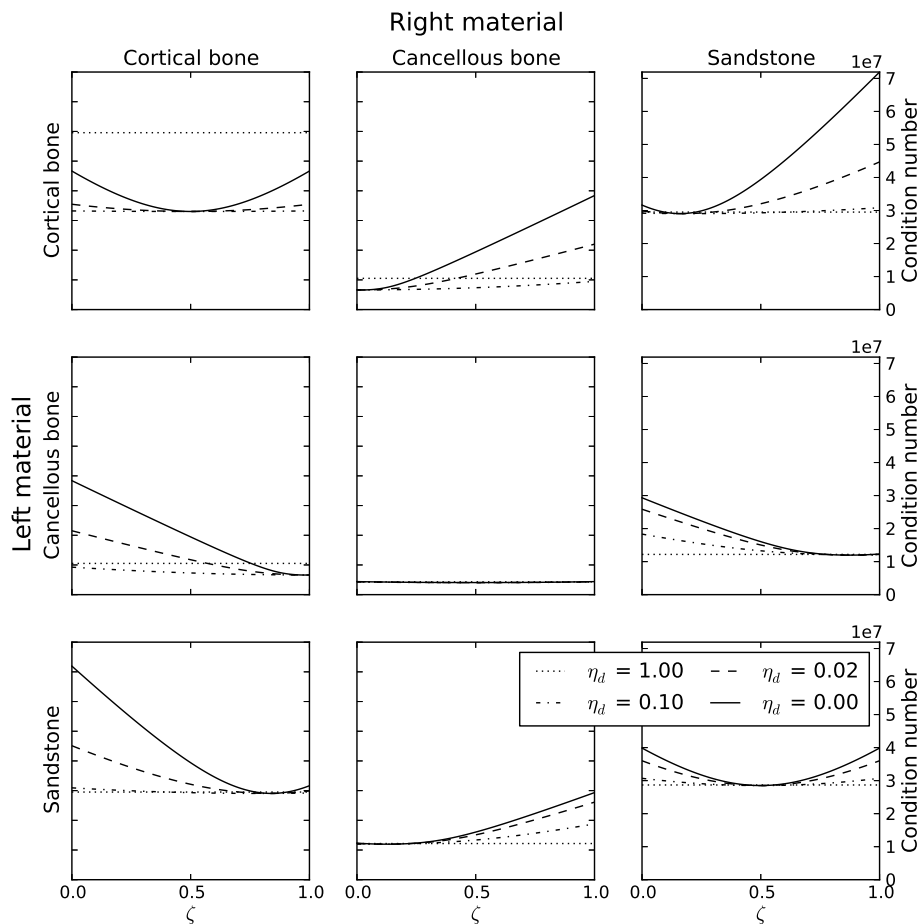


Figure 5.2: Variation of the condition number of the Riemann solution linear system (5.14) for interfaces between various pairs of poroelastic materials, as a function of the free parameter ζ . Material properties are given in Table A.1. The maximum condition number observed is around 7×10^7 — not large enough to endanger accuracy in IEEE double-precision arithmetic.

essentially irrelevant except for impermeable or nearly impermeable interfaces, and even for these interfaces there is no worse than about a factor of five difference between the lowest and highest condition number for a given material combination — not a significant variation, given that the highest condition number observed is around 7×10^7 . For similar materials, the minimum condition number occurs near $\zeta = \frac{1}{2}$, and even for very dissimilar materials this value of ζ yields a condition number substantially below the worst possible. Based on these results, for simplicity this free parameter will be set to $\zeta = \frac{1}{2}$ for all further computations, corresponding to equal weighting of both sides in computing $\widehat{\mathbf{q}} \cdot \mathbf{n}$.

5.3 Modifications to the high-resolution finite volume scheme

The presence of interfaces between media governed by different PDEs presents a special problem for implementing certain aspects of high-resolution finite volume methods. To explore this in more detail, we first write the full update formula for cell ij from time step n to $n + 1$ on a mapped grid, following (21.12) and (23.11) from [57]:

$$\begin{aligned} \mathbf{Q}_{ij}^{n+1} = & \mathbf{Q}_{ij} \\ & - \frac{\Delta t}{\kappa_{ij} \Delta \xi_1} \left(\mathcal{A}^+ \Delta \mathbf{Q}_{i-1/2,j} + \mathcal{A}^- \Delta \mathbf{Q}_{i+1/2,j} \right) - \frac{\Delta t}{\kappa_{ij} \Delta \xi_2} \left(\mathcal{B}^+ \Delta \mathbf{Q}_{i,j-1/2} + \mathcal{B}^- \Delta \mathbf{Q}_{i,j+1/2} \right) \\ & - \frac{\Delta t}{\kappa_{ij} \Delta \xi_1} \left(\tilde{\mathbf{F}}_{i+1/2,j} - \tilde{\mathbf{F}}_{i-1/2,j} \right) - \frac{\Delta t}{\kappa_{ij} \Delta \xi_2} \left(\tilde{\mathbf{G}}_{i,j+1/2} - \tilde{\mathbf{G}}_{i,j-1/2} \right). \end{aligned} \quad (5.19)$$

All quantities on the right-hand side are evaluated at time step n . The coordinates ξ_1 and ξ_2 are computational coordinates, κ_{ij} is the ratio of the area of cell ij to $\Delta \xi_1 \Delta \xi_2$, $\mathcal{A}^+ \Delta \mathbf{Q}_{i-1/2,j}$ etc. are the fluctuations computed from the Riemann solutions at the interfaces with neighboring cells, and $\tilde{\mathbf{F}}_{i-1/2,j}$ etc. are the correction fluxes. These correction fluxes are a combination of corrections coming from transverse Riemann solutions, which are important for stability in multi-dimensional problems, and higher-order correction fluxes, which allow the method to achieve second-order accuracy.

It is these correction fluxes that pose a problem here. Because at a fluid-poroelastic interface the solutions in the two domains are governed by different PDEs, it is not possible to construct a single flux transferring quantities from one domain to the other in this form. For example, in the formulation of acoustics used here, elements 2 through 6 of \mathbf{Q} are identically zero in a fluid, because they correspond to the total stress tensor and solid velocity. The state of stress in a fluid is described completely by the pressure (since viscosity is ignored in an all-fluid medium for this work), and there is no solid component present to have a velocity, so a flux appropriate to the poroelastic medium would produce nonsensical results if used to update the solution in the fluid. Similarly, a flux appropriate to the fluid would only update the pressure and fluid flow rate in the poroelastic medium, whereas we should in general expect all of the state variables to be updated. The update formula (5.19) must therefore be reformulated to correctly handle a fluid-poroelastic interface.

Because a single flux cannot be defined across an interface between media of different types, (5.19) will instead be reformulated with two correction fluctuations, one on either side of the interface — instead of a correction flux $\tilde{\mathbf{G}}_{i,j+1/2}$ acting on both cells (i, j) and $(i, j + 1)$, define a $\tilde{\mathbf{G}}_{i,j+1/2}^+$ acting on cell $(i, j + 1)$ and a $\tilde{\mathbf{G}}_{i,j+1/2}^-$ acting on cell (i, j) . These fluctuations have a similar meaning to the first-order fluctuations $\mathcal{A}^\pm \Delta \mathbf{Q}$ — they measure the rate of change of cell averages caused by waves propagating from the interface. The correction fluctuations are not necessarily equal on either side of the interface, but they are compatible in the sense that they arise from the same Riemann problems and respect the underlying physics of both media. In terms of the correction fluctuations, the new update

formula derived from (5.19) is

$$\begin{aligned} \mathbf{Q}_{ij}^{n+1} = & \mathbf{Q}_{ij} \\ & - \frac{\Delta t}{\kappa_{ij} \Delta \xi_1} \left(\mathcal{A}^+ \Delta \mathbf{Q}_{i-1/2,j} + \mathcal{A}^- \Delta \mathbf{Q}_{i+1/2,j} \right) - \frac{\Delta t}{\kappa_{ij} \Delta \xi_2} \left(\mathcal{B}^+ \Delta \mathbf{Q}_{i,j-1/2} + \mathcal{B}^- \Delta \mathbf{Q}_{i,j+1/2} \right) \\ & - \frac{\Delta t}{\kappa_{ij} \Delta \xi_1} \left(\tilde{\mathbf{F}}_{i+1/2,j}^- - \tilde{\mathbf{F}}_{i-1/2,j}^+ \right) - \frac{\Delta t}{\kappa_{ij} \Delta \xi_2} \left(\tilde{\mathbf{G}}_{i,j+1/2}^- - \tilde{\mathbf{G}}_{i,j-1/2}^+ \right). \end{aligned} \quad (5.20)$$

Note that the new correction fluctuations keep the same sign convention as the old correction fluxes. Also note that for a Riemann problem with an interface condition of the form (5.12), the solution is not necessarily a function purely of the difference in cell states; the $\mathcal{A}^\pm \Delta \mathbf{Q}_{i-1/2,j}$ notation is retained for familiarity's sake, but it should be interpreted as fluctuations arising from a Riemann problem at cell interface $(i-1/2, j)$, not as an operator applied to a difference of states.

The following subsections describe how the correction fluctuations are computed.

5.3.1 Transverse Riemann solution

The transverse Riemann solution process computes the contribution of the solution on a cell in one time step to solutions on the cells diagonally adjacent to it in the next time step, and is important for stability in a dimensionally-unsplit high-resolution finite volume scheme. Consider the contribution to \mathbf{Q}_{ij}^{n+1} of the normal Riemann solution at left edge of cell (i, j) . Ordinarily, the transverse Riemann solution process would use the right-going fluctuation $\mathcal{A}^+ \Delta \mathbf{Q}_{i-1/2,j}$ to calculate up-going and down-going transverse fluctuations $\mathcal{B}^+ \mathcal{A}^+ \Delta \mathbf{Q}_{i-1/2,j}$ and $\mathcal{B}^- \mathcal{A}^+ \Delta \mathbf{Q}_{i-1/2,j}$. This is illustrated in Figure 5.3a. In the case of a linear system such as this, these transverse fluctuations are found by decomposing $\mathcal{A}^+ \Delta \mathbf{Q}_{i-1/2,j}$ into eigenvectors of the flux Jacobians corresponding to the interfaces between cell (i, j) and cells $(i, j-1)$ and $(i, j+1)$, then multiplying by the respective eigenvalues and by the ratio of the physical length of the cell interface to $\Delta \xi_1$; the portion of the eigendecomposition corresponding to positive eigenvalues is $\mathcal{B}^+ \mathcal{A}^+ \Delta \mathbf{Q}_{i-1/2,j}$, while the portion corresponding to negative eigenvalues is $\mathcal{B}^- \mathcal{A}^+ \Delta \mathbf{Q}_{i-1/2,j}$. The correction fluxes $\tilde{\mathbf{G}}_{i,j+1/2}$ and $\tilde{\mathbf{G}}_{i,j-1/2}$ are then incremented by $-\frac{\Delta t}{2\Delta \xi_1} \mathcal{B}^+ \mathcal{A}^+ \Delta \mathbf{Q}_{i-1/2,j}$ and $-\frac{\Delta t}{2\Delta \xi_1} \mathcal{B}^- \mathcal{A}^+ \Delta \mathbf{Q}_{i-1/2,j}$, respectively.

For the new approach, illustrated in Figure 5.3b, the contributions of the fluctuation $\mathcal{A}^+ \Delta \mathbf{Q}_{i-1/2,j}$ to, for instance, the two distinct correction fluctuations $\tilde{\mathbf{G}}_{i,j-1/2}^\pm$ are computed by applying the Riemann solver to the states in cells (i, j) and $(i, j-1)$ modified by the fluctuations from the normal Riemann solve. (For the poroelasticity and acoustics systems considered here, the solver exploits linearity by separately computing the transverse fluctuations from each normal fluctuation, which is the easiest approach to fit into the existing CLAWPACK framework. It would also be possible to reduce computational effort by computing combined transverse fluctuations from the normal fluctuations in cells in two adjacent rows taken together — e.g., computing $\mathcal{A}^+ \Delta \mathbf{Q}_{i-1/2,j}$ and $\mathcal{A}^+ \Delta \mathbf{Q}_{i-1/2,j+1}$, then using them together to compute $\mathcal{B}_{i,j+1/2}^\pm$ fluctuations by a process very similar to solving a normal Riemann problem between states $\mathcal{A}^+ \Delta \mathbf{Q}_{i-1/2,j}$ and $\mathcal{A}^+ \Delta \mathbf{Q}_{i-1/2,j+1}$.) The fluctuations $\mathcal{B}^\pm \mathcal{A}^+ \Delta \mathbf{Q}_{i-1/2,j}$ from the Riemann solution are then used to increment $\tilde{\mathbf{G}}_{i,j-1/2}^-$




Figure 5.3: Comparison of classical and fluctuation-based approaches to cell updates from transverse Riemann solves. Classical approach: both cells on either side of the transverse interface are updated using the same transverse flux. Fluctuation-based approach: each cell is updated using a separate fluctuation.

by $-\frac{\Delta t}{2\Delta\xi_1}\mathcal{B}^-\mathcal{A}^+\Delta\mathbf{Q}_{i-1/2,j}$, and $\tilde{\mathbf{G}}_{i,j-1/2}^+$ by $+\frac{\Delta t}{2\Delta\xi_1}\mathcal{B}^+\mathcal{A}^+\Delta\mathbf{Q}_{i-1/2,j}$. The other correction fluctuations are updated analogously from transverse Riemann solutions computed at each cell interface.

To compare these two approaches, let us examine the full contribution to the update of cell (i, j) of the transverse solves based on the right-going fluctuation from its left edge. To make the difference between the methods clear, a subscript will be attached to each transverse fluctuation operator identifying the interface at which it was computed — so, for example, the up-going transverse fluctuation from the upper edge of cell (i, j) is $\mathcal{B}_{i,j+1/2}^+\mathcal{A}^+\Delta\mathbf{Q}_{i-1/2,j}$. The classical transverse solve approach increments cell (i, j) by

$$\Delta\mathbf{Q}_{ij,\text{trans,classical}} = -\frac{\Delta t^2}{2\kappa_{ij}\Delta\xi_1\Delta\xi_2} \left(-\mathcal{B}_{i,j+1/2}^+\mathcal{A}^+\Delta\mathbf{Q}_{i-1/2,j} + \mathcal{B}_{i,j-1/2}^-\mathcal{A}^+\Delta\mathbf{Q}_{i-1/2,j} \right), \quad (5.21)$$

while the new fluctuation-oriented approach increments the cell by

$$\Delta\mathbf{Q}_{ij,\text{trans,fluctuation}} = -\frac{\Delta t^2}{2\kappa_{ij}\Delta\xi_1\Delta\xi_2} \left(\mathcal{B}_{i,j+1/2}^-\mathcal{A}^+\Delta\mathbf{Q}_{i-1/2,j} - \mathcal{B}_{i,j-1/2}^+\mathcal{A}^+\Delta\mathbf{Q}_{i-1/2,j} \right). \quad (5.22)$$

The two approaches produce identical results if the $\mathcal{B}_{i,j+1/2}^\pm$ operators are the same as $\mathcal{B}_{i,j-1/2}^\pm$, but differ otherwise. Notice that (5.22) updates the cell using fluctuations computed from waves entering it, while (5.21) uses fluctuations from waves leaving it, which may not be appropriate. Both approaches produce the same increment to cell ij from the transverse solves applied to the normal fluctuations from the rows above and below, $\mathcal{A}^+ \Delta \mathbf{Q}_{i-1/2,j\pm 1}$.

One final note should be made on the broader applicability of this way of handling transverse propagation, regarding conservation. While the governing equations used here for a fluid-poroelastic system are not cast in conservation form, one of the great strengths of high-resolution finite volume methods is their ability to maintain conservation of physically relevant quantities such as mass, momentum, or energy when they are applied to systems of conservation laws. Going from a single transverse flux to a pair of transverse fluctuations appears to jeopardize this — it is not obvious that $\mathcal{B}_{i,j+1/2}^+ \mathcal{A}_{i-1/2,j}^+ \Delta \mathbf{Q}$ will update cell $(i, j+1)$ by an equal and opposite amount to the update of $\mathcal{B}_{i,j+1/2}^- \mathcal{A}_{i-1/2,j}^+ \Delta \mathbf{Q}$ to cell (i, j) . Ensuring that the transverse solve maintains conservation, though, is very similar to ensuring that the two normal fluctuations $\mathcal{B}^\pm \Delta \mathbf{Q}$ update neighboring cells in a conservative fashion, so it should be no barrier to practical use. To demonstrate, suppose that we have some general, possibly nonlinear problem. Let $\mathcal{B}^\pm(\mathbf{Q}_{i,j+1}, \mathbf{Q}_{ij})$ denote the up-going and down-going fluctuations from the normal Riemann solver applied to states $\mathbf{Q}_{i,j+1}$ and \mathbf{Q}_{ij} in cells $(i, j+1)$ and (i, j) . Assume that for all i and j these fluctuations satisfy the condition

$$\mathcal{B}^+(\mathbf{Q}_{i,j+1}, \mathbf{Q}_{ij}) + \mathcal{B}^-(\mathbf{Q}_{i,j+1}, \mathbf{Q}_{ij}) = \mathbf{G}(\mathbf{Q}_{i,j+1}) - \mathbf{G}(\mathbf{Q}_{ij}). \quad (5.23)$$

This condition is sufficient to ensure conservation [57]; the function \mathbf{G} will typically be the vertical component of the flux vector of the underlying conservation law, but this is not necessary. In that case, a sufficient condition for the new transverse solution scheme to maintain conservation is to require the output of the transverse Riemann solve to be related to that of the normal Riemann solve by

$$\begin{aligned} & \mathcal{B}_{i,j+1/2}^\pm \mathcal{A}_{i-1/2,j}^+ \Delta \mathbf{Q} + \mathcal{B}_{i,j+1/2}^\pm \mathcal{A}_{i-1/2,j+1}^+ \Delta \mathbf{Q} \\ &= \mathcal{B}^\pm(\mathbf{Q}_{i,j+1} + \mathcal{A}_{i-1/2,j+1}^+ \Delta \mathbf{Q}, \mathbf{Q}_{ij} + \mathcal{A}_{i-1/2,j}^+ \Delta \mathbf{Q}) - \mathcal{B}^\pm(\mathbf{Q}_{i,j+1}, \mathbf{Q}_{ij}). \end{aligned} \quad (5.24)$$

Essentially, this condition states that the total of the transverse fluctuations should be equal to the difference between the up- and down-going normal fluctuations from the states behind the right-going waves, and from the states in front of them.

If condition (5.24) holds, then summing the transverse fluctuations at interface $(i, j+1/2)$ and rearranging a bit gives

$$\begin{aligned} & \mathcal{B}_{i,j+1/2}^+ \mathcal{A}_{i-1/2,j}^+ \Delta \mathbf{Q} + \mathcal{B}_{i,j+1/2}^+ \mathcal{A}_{i-1/2,j+1}^+ \Delta \mathbf{Q} + \mathcal{B}_{i,j+1/2}^- \mathcal{A}_{i-1/2,j}^+ \Delta \mathbf{Q} + \mathcal{B}_{i,j+1/2}^- \mathcal{A}_{i-1/2,j+1}^+ \Delta \mathbf{Q} \\ &= (\mathbf{G}(\mathbf{Q}_{i,j+1} + \mathcal{A}_{i-1/2,j+1}^+ \Delta \mathbf{Q}) - \mathbf{G}(\mathbf{Q}_{i,j+1})) - (\mathbf{G}(\mathbf{Q}_{ij} + \mathcal{A}_{i-1/2,j}^+ \Delta \mathbf{Q}) - \mathbf{G}(\mathbf{Q}_{ij})) \\ &= \mathbf{K}_{i,j+1} - \mathbf{K}_{i,j}, \end{aligned} \quad (5.25)$$

where $\mathbf{K}_{ij} := \mathbf{G}(\mathbf{Q}_{ij} + \mathcal{A}_{i-1/2,j}^+ \Delta \mathbf{Q}) - \mathbf{G}(\mathbf{Q}_{ij})$. This is precisely analogous to (5.23), and ensures that the transverse fluctuations will update the solution in a conservative fashion.

Condition (5.24) is typically satisfied by transverse solvers for linear problems (including the transverse solvers used in this work away from interfaces between distinct media, which is the only situation where conservation of the elements of \mathbf{Q} makes sense for the problems considered here), and (5.24) may be taken as a recipe for defining conservative transverse fluctuations for problems for which it is not immediately obvious how to do so.

5.3.2 Second-order correction term

These terms modify the basic Godunov-type update of the first line of (5.19) with an additional antidiffusive flux. Mathematically, they provide the second-order terms in the Taylor expansion of the solution in time; qualitatively, they remove the diffusive error from the Godunov update, sharpening the solution. In a classical high-resolution finite volume method, these second-order correction fluxes are computed as linear combinations of the waves propagating from each interface, with *limiters* applied to each wave. These limiters compare the strength of a wave at an interface to that of the wave in the same family at the neighboring interface in the upwind direction; based on its relative strength the wave may be scaled downward in magnitude to avoid overshoot, or amplified to sharpen the solution. For a more thorough discussion, see Chapter 6 of LeVeque [57].

There are two problems with implementing appropriate second-order correction fluxes here. The first is in formulating the appropriate numerical flux function on either side of an interface where a condition of the form (5.12) holds. This is not trivial, but can be accomplished with, for instance, an approach based on the Immersed Interface Method [59, 72, 73]. The second, harder problem is in creating an appropriate limiting scheme when the waves in the upwind direction are in a different medium and may be in no way analogous to the waves we wish to limit.

Because of these problems, the second-order correction term is omitted at interfaces between different materials. While this reduces the accuracy of the solution to first order locally, note that classical high-resolution finite volume methods also lose formal accuracy at such interfaces, even if the correction term is included (Section 9.12 of [57]), and that it is only at interfaces between different materials that higher error is incurred. At cell interfaces between identical materials, the standard second-order correction flux is used.

5.4 Source term and numerical software

As in the previous chapters, the source term \mathbf{DQ} in (2.43) is handled using operator splitting. Based on the results of Chapter 4, only Strang splitting is used here. The exact solution operator $\exp(\mathbf{D}\Delta t)$ is again used to advance the solution under the action of the dissipation term.

The numerical solution techniques described in this chapter are again implemented using the CLAWPACK finite volume method package, version 4.6 [71]. Normally, writing a CLAWPACK application would require only writing a few plug-in subroutines, such as the Riemann solver, but because of the more extensive modifications to the standard high-resolution finite volume formulation, some of its internal subroutines also had to be modified. Using CLAWPACK still greatly reduced the time and effort required for coding compared to starting from scratch, however. The new transverse propagation scheme of Section 5.3.1 was

implemented in AMRCLAW because the structure of the AMR code required fewer invasive changes than standard CLAWPACK; however, no adaptive mesh refinement was used in the two-dimensional results following in chapters 6 and 7.

Chapter 6

**RESULTS FOR POROELASTIC-FLUID SYSTEMS ON
TWO-DIMENSIONAL RECTILINEAR GRIDS**

This chapter tests the ability of the code implementing the methods of Chapter 5 to correctly model interfaces between different poroelastic materials, and between a poroelastic material and a fluid, on a rectilinear grid. Note that although the grid lines are straight, these results still use the mapped-grid solver; the grid mapping just happens to be particularly simple. For all cases in this chapter, the simulation time step is chosen so that the global maximum CFL number is 0.9. No limiters are used anywhere in these simulations — while limiters improve solution accuracy on typical grids, they obscure the convergence behavior of the underlying wave propagation method, which is what is being investigated here. Since the solutions used in this chapter are well-resolved everywhere, and smooth except at interfaces between different media, the numerical solution should not encounter serious trouble from the numerical dispersion that limiters are designed to suppress. Table A.1 lists the properties of all the materials used here.

The results shown in this chapter are error values and convergence rates. Error is measured relative to the true solution using grid energy norms: the energy norm defined by (2.54) is taken of the difference between the numerical solution for each grid cell and the true solution evaluated at the cell centroid. The 1-norm and max-norm errors are then defined as

$$\text{error 1-norm} = \frac{1}{N_1 N_2} \sum_{i=1}^{N_1} \sum_{j=1}^{N_2} \|\mathbf{Q}_{ij,\text{numerical}} - \mathbf{Q}_{ij,\text{true}}\|_E \quad (6.1)$$

$$\text{error max-norm} = \max_{i,j} \|\mathbf{Q}_{ij,\text{numerical}} - \mathbf{Q}_{ij,\text{true}}\|_E$$

where N_1 and N_2 are the grid dimensions in the two computational axes. The incident waves are all scaled to unit peak energy density, so these grid energy error norms provide an overall measure of relative error in the numerical solution, with the various components of the solution scaled appropriately relative to each other. All cases in this chapter test the ability of the code to correctly evolve a known analytical solution: all simulations are initialized by evaluating the true solution at cell centroids, and boundary conditions are implemented using the standard ghost cell approach, with ghost cells filled using the true solution evaluated at ghost cell centroids. In addition, all the solutions considered in this chapter are periodic in time. In such cases the error in the numerical solution displays a periodic component, which gives it local minima at integer multiples of the period after the starting time, so to avoid artificially picking the best result the solution error is evaluated after 1.25 periods. Since a very large number of simulations are run, convergence is not generally displayed using plots of error versus grid size or spacing. Instead, for each true solution a linear least-squares fit of the logarithm of error in each norm with respect to the logarithm of grid size is performed; the best, worst, and mean convergence rates are

Table 6.1: Convergence results for acoustic and inviscid poroelastic waves with the second-order correction term omitted along a line of cell interfaces.

	Error norm	Convergence rate			Worst R^2 value	Error on 800×800 grid	
		Best	Worst	Mean		Best	Worst
Acoustic	1-norm	2.01	1.99	2.00	0.99957	2.57×10^{-5}	4.09×10^{-5}
	Max-norm	1.49	1.03	1.15	0.95837	2.76×10^{-4}	2.51×10^{-3}
Fast P	1-norm	2.02	1.97	1.99	0.99883	2.75×10^{-5}	7.83×10^{-5}
	Max-norm	1.45	0.91	1.13	0.86779	4.62×10^{-4}	3.74×10^{-3}
S	1-norm	2.01	1.98	1.99	0.99999	4.18×10^{-5}	9.49×10^{-5}
	Max-norm	1.76	0.95	1.15	0.96214	4.05×10^{-4}	5.06×10^{-3}
Slow P	1-norm	2.01	1.95	1.98	0.99999	8.21×10^{-5}	3.10×10^{-4}
	Max-norm	1.85	0.76	1.09	0.98524	9.91×10^{-4}	1.99×10^{-2}

reported across all solutions, as measured by the slopes of these fit lines, along with the worst R^2 value for the fit, and the best and worst errors in each norm on the finest grid used. For all test cases in this chapter, convergence is assessed using grid dimensions of 100×100 , 200×200 , 400×400 , and 800×800 cells.

Before treating material interfaces on rectilinear grids, let us first explore the consequences of omitting the second-order correction term along a line of cell interfaces.

6.1 Effect of omitting the second-order correction term

The test cases used to observe the results of omitting the second-order correction term are time-harmonic simple plane waves propagating in homogeneous media. The correction term is omitted for the vertical-direction fluxes along one grid line passing horizontally through the center of the domain. The results are examined both for poroelastic waves of all three families — though with viscosity omitted, to isolate the effect of the missing term — and for acoustic waves; the material used is the orthotropic sandstone of Table A.1 for the poroelastic waves, or just the brine contained within it for acoustic waves. To test grid alignment and transverse solve effects, the wave propagation direction varies from straight down in the $-z$ direction to 7.5° clockwise of the $+x$ direction, in 7.5° increments. Because the sample poroelastic material is orthotropic, its principal 1-direction is also varied from horizontal to 165° counterclockwise of horizontal, in 15° increments. Since these tests are inviscid, the period of the wave and time required to cross the domain are the only time scales present, so an angular frequency of 1 rad/s is used more or less arbitrarily. For all cases, the computational domain is a square whose side length is two wavelengths of a wave in the incident family propagating in the material principal 1-direction.

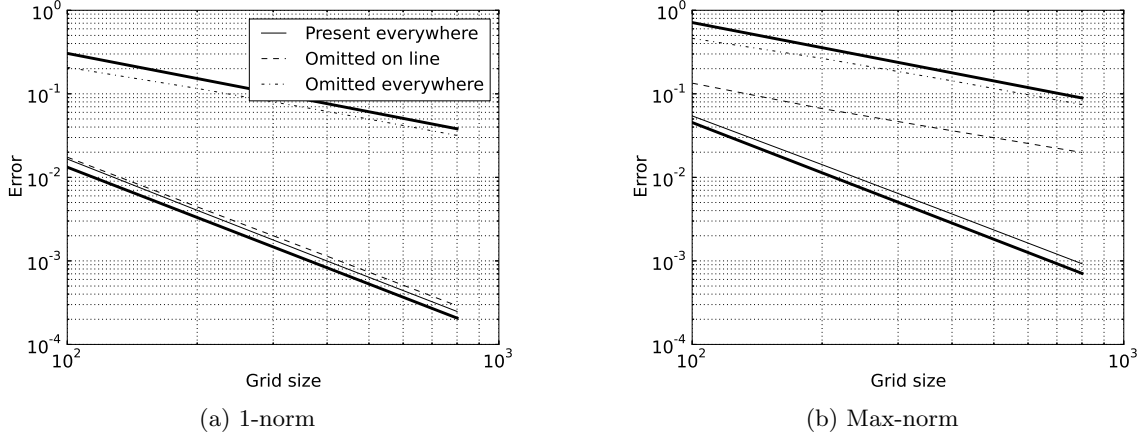


Figure 6.1: Detailed examination of the effect of omitting the second-order correction term for the case of Table 6.1 that had the highest error on the finest grid. Results of omitting the second-order correction term on a line are compared to results if the term is present or omitted everywhere. The thick black lines are first- and second-order reference lines.

Table 6.1 shows the results of this convergence study. Even with the second-order correction term omitted along a line, the code still achieves second-order convergence in the 1-norm. Convergence degrades to first-order (typically) or somewhat below (in the worst case) for the max-norm. These results are about as good as could be expected from omitting the second-order term. Figure 6.1 examines in greater detail the case that gives the worst error on the 800×800 grid — a slow P wave propagating vertically downward, with the principal direction of the poroelastic sandstone medium oriented 15° counterclockwise of horizontal. In the 1-norm, results are almost as good as if the second-order correction term were present everywhere, while in the max-norm, even though the convergence rate is reduced to first-order when the term is omitted along a line, the magnitude of the error is almost four times less than if it were omitted everywhere. This suggests that we can expect results of reasonable accuracy even with the second-order correction term omitted at interfaces between different media.

6.2 Reflected and transmitted waves at a fluid-poroelastic interface

To test the code’s ability to correctly handle material interfaces, let us first examine the case of a time-harmonic plane acoustic wave striking a flat interface between a poroelastic medium and a fluid. In all cases, the interface between the two media is horizontal, and the incident wave impinges from the top side. Figure 6.2a shows a simple sketch of the problem. Analytical solutions are generated for these cases using the procedure described in Appendix B. The poroelastic medium used is the brine-saturated sandstone of Table A.1, and the fluid medium is the same brine contained in the sandstone. A combinatorial

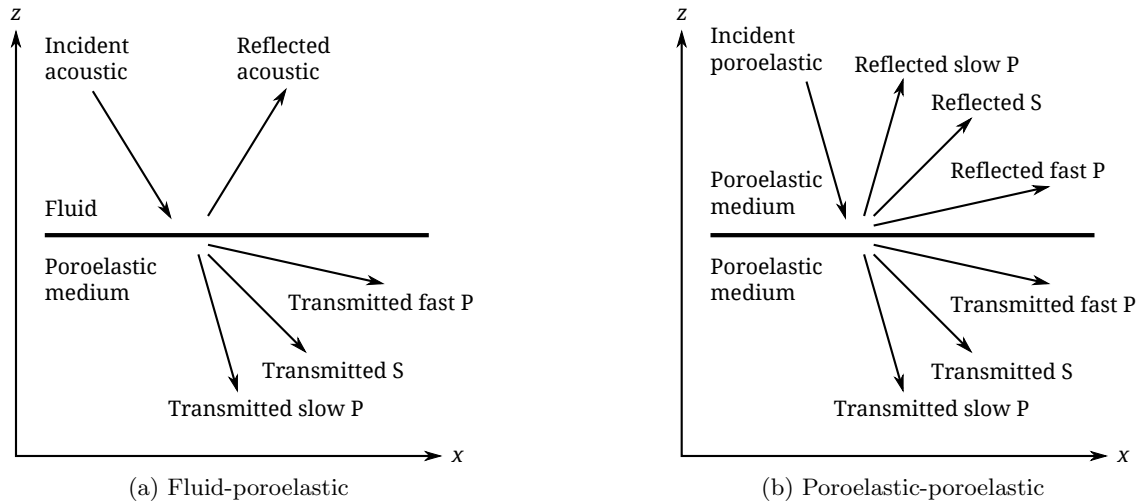


Figure 6.2: Simple sketches of reflection/transmission problems used to test convergence. The incident, reflected, and transmitted waves are labeled by type; an incident poroelastic wave may be of any of the three types.

sweep is again performed over the relevant parameters. In this case, the direction of the incident wave is varied from 7.5° below horizontal to straight down in 7.5° increments, and the principal material 1-direction is varied from horizontal to 165° counterclockwise of horizontal in 15° increments. For each combination of incident wave and material principal directions, simulations are conducted with interface discharge efficiencies η_d of 0 (sealed pores), 0.5 (imperfect hydraulic contact), and 1 (open pores) used in interface condition (2.105). The dimensions of the domain are two wavelengths of an acoustic wave at the chosen frequency. These tests are performed both with viscosity ignored, in order to investigate the performance of the hyperbolic solver by itself, and with viscosity included. For the inviscid tests, since the period of the wave and time required for it to cross the domain are the only time scales present, the wave frequency is not directly relevant to the solution error, so an angular frequency of 1 rad/s is used. With viscosity included, the dissipation term has its own intrinsic time scale independent of the wave behavior, so choice of frequency is important. These viscous tests are restricted to a high frequency (10 kHz; the maximum frequency for validity of Biot theory in the sandstone medium is 25 kHz), because at low frequencies the transmitted slow P wave dissipates over distances much shorter than a wavelength. In fact, the characteristic decay length is much shorter than the grid spacing on an otherwise reasonably resolved grid, so at low frequencies it would not be possible to resolve the transmitted slow P wave without refining the grid to impractical levels. This also results in grid sizes and corresponding time steps short enough that these models are outside the stiff regime identified in Chapter 4. For all cases, the computational domain is again a square with side length equal to two wavelengths of incident acoustic wave.

Table 6.2: Convergence results for an acoustic wave impinging on sandstone, with viscosity omitted.

	Error norm	Convergence rate			Worst R^2 value	Error on 800×800 grid	
		Best	Worst	Mean		Best	Worst
$\eta_d = 0$	1-norm	2.00	1.98	1.99	1.00000	7.93×10^{-5}	2.09×10^{-4}
	Max-norm	1.50	0.88	1.12	0.96350	2.17×10^{-3}	1.25×10^{-2}
$\eta_d = 0.5$	1-norm	2.01	1.97	1.99	0.99999	1.03×10^{-4}	2.94×10^{-4}
	Max-norm	1.47	0.88	1.14	0.95713	2.12×10^{-3}	1.24×10^{-2}
$\eta_d = 1$	1-norm	2.01	1.97	2.00	0.99999	1.07×10^{-4}	3.14×10^{-4}
	Max-norm	1.52	0.88	1.15	0.95981	2.16×10^{-3}	1.24×10^{-2}

Table 6.3: Convergence results for a 10 kHz acoustic wave impinging on sandstone, with viscosity included.

	Error norm	Convergence rate			Worst R^2 value	Error on 800×800 grid	
		Best	Worst	Mean		Best	Worst
$\eta_d = 0$	1-norm	1.99	1.96	1.98	1.00000	6.93×10^{-5}	1.21×10^{-4}
	Max-norm	1.44	0.88	1.09	0.96397	2.46×10^{-3}	1.25×10^{-2}
$\eta_d = 0.5$	1-norm	1.97	1.90	1.95	0.99999	6.84×10^{-5}	1.79×10^{-4}
	Max-norm	1.30	0.72	0.98	0.97612	2.96×10^{-3}	2.06×10^{-2}
$\eta_d = 1$	1-norm	1.97	1.90	1.94	0.99999	6.75×10^{-5}	1.83×10^{-4}
	Max-norm	1.29	0.70	0.98	0.97672	3.02×10^{-3}	2.19×10^{-2}

Tables 6.2 and 6.3 show the results of these convergence studies for the inviscid and viscous high-frequency tests respectively. Second-order convergence is obtained for all cases in the 1-norm, though results are degraded to first-order in the max-norm due to the omission of the second-order correction term at the interface, as expected from the results of the previous section. All three interface conditions (open and sealed pores, and imperfect hydraulic contact) show similar levels of error, and the difference between best and worst error obtained in the finest grid is also not large. Convergence rates are slightly worse for the viscous cases; they are still typically quite close to the expected second-order and first-order results, but some outlying cases degrade as far as order 0.7, possibly due to the compounding of first-order error at the interface with operator splitting error.

6.3 *Reflected and transmitted waves at an interface between two poroelastic materials*

Next, let us examine plane waves reflected and transmitted at an interface between two poroelastic media. For these cases the upper medium is the isotropic, brine-saturated shale of Table A.1, and the lower medium is again orthotropic sandstone; Figure 6.2b shows a simple sketch of the reflected and transmitted waves in this case. These cases are quite similar to the fluid-poroelastic cases, with the exception that incident waves in multiple different poroelastic wave families are sent in. To reduce the total number of cases, a smaller number of incident wave directions are used here — the angle of incidence varies from 7.5° below horizontal to 82.5° in steps of 15° , which halves the number of angles of incidence and drops the relatively uninteresting normal incidence, while retaining the grazing 7.5° angle. The same set of principal material directions for the sandstone is used as in the fluid-poroelastic cases; principal material direction is irrelevant for the shale because it is isotropic. The same sets of interface discharge efficiencies and grid dimensions are also used. For all cases the domain size is two wavelengths of the fast P wave in shale at the chosen frequency. The inviscid cases use waves with an angular frequency of 1 rad/s, while the viscous cases use 20 kHz waves, a frequency high enough to be out of the stiff regime (or just at the edge of it in shale for a 100×100 grid) but low enough for low-frequency Biot theory to be valid for both materials. Incident waves in all three families are simulated for the inviscid cases in order to exercise all different possibilities in the solution code, but with viscosity present, incident slow P waves are not included — their decay rate is high enough that their amplitude is reduced by a factor of 10^{20} over the width of the domain, making these cases both intractable to simulate and unrealistic.

Tables 6.4 and 6.5 show the results of these convergence studies. Similarly to the acoustic-poroelastic cases, consistent second-order convergence occurs in the 1-norm, with first-order or better convergence in the max-norm. This occurs for all incident wave types, even in the face of the operator splitting error present in the viscous cases. In fact, the worst-case convergence rates are better than for the acoustic-poroelastic cases, perhaps due to a more tractable interface condition.

Table 6.4: Convergence results for poroelastic waves in shale impinging on sandstone, with viscosity omitted.

	Incident wave	Error norm	Convergence rate			Worst R^2 value	Error on 800×800 grid	
			Best	Worst	Mean		Best	Worst
$\eta_d = 0$	Fast P	1-norm	2.01	1.99	2.00	1.00000	8.81×10^{-5}	3.26×10^{-4}
		Max-norm	1.49	1.06	1.17	0.97570	2.01×10^{-3}	5.62×10^{-3}
	S	1-norm	2.01	1.98	1.99	0.99999	2.19×10^{-4}	4.27×10^{-4}
		Max-norm	1.86	1.03	1.25	0.98292	2.39×10^{-3}	1.18×10^{-2}
	Slow P	1-norm	2.00	1.99	2.00	0.99999	3.34×10^{-4}	6.34×10^{-4}
		Max-norm	1.92	1.08	1.40	0.99555	3.90×10^{-3}	2.13×10^{-2}
$\eta_d = 0.5$	Fast P	1-norm	2.00	1.98	1.99	1.00000	7.65×10^{-5}	2.00×10^{-4}
		Max-norm	1.22	1.01	1.13	0.99179	1.99×10^{-3}	5.63×10^{-3}
	S	1-norm	2.01	1.98	1.99	0.99999	2.13×10^{-4}	3.49×10^{-4}
		Max-norm	1.86	0.99	1.24	0.98301	2.38×10^{-3}	1.18×10^{-2}
	Slow P	1-norm	2.01	1.99	2.00	0.99999	4.70×10^{-4}	1.44×10^{-3}
		Max-norm	1.99	1.06	1.40	0.96976	3.81×10^{-3}	1.96×10^{-2}
$\eta_d = 1$	Fast P	1-norm	2.00	1.99	1.99	1.00000	7.66×10^{-5}	1.99×10^{-4}
		Max-norm	1.22	1.00	1.13	0.99181	1.99×10^{-3}	5.63×10^{-3}
	S	1-norm	2.01	1.98	1.99	0.99999	2.11×10^{-4}	3.49×10^{-4}
		Max-norm	1.86	0.98	1.24	0.98301	2.38×10^{-3}	1.18×10^{-2}
	Slow P	1-norm	2.01	1.99	2.00	0.99999	4.86×10^{-4}	1.50×10^{-3}
		Max-norm	1.99	1.05	1.41	0.97068	3.86×10^{-3}	1.98×10^{-2}

Table 6.5: Convergence results for 20 kHz poroelastic waves in shale impinging on sandstone, with viscosity included.

	Incident wave	Error norm	Convergence rate			Worst R^2 value	Error on 800×800 grid	
			Best	Worst	Mean		Best	Worst
$\eta_d = 0$	Fast P	1-norm	2.00	1.96	1.99	1.00000	6.97×10^{-5}	1.80×10^{-4}
		Max-norm	1.21	1.01	1.12	0.99164	1.98×10^{-3}	5.88×10^{-3}
	S	1-norm	2.01	1.98	1.99	1.00000	2.10×10^{-4}	4.40×10^{-4}
		Max-norm	1.86	1.02	1.25	0.97550	2.56×10^{-3}	1.20×10^{-2}
$\eta_d = 0.5$	Fast P	1-norm	2.00	1.97	1.99	1.00000	6.99×10^{-5}	1.81×10^{-4}
		Max-norm	1.21	1.01	1.12	0.99168	1.98×10^{-3}	5.78×10^{-3}
	S	1-norm	2.01	1.99	1.99	0.99999	2.09×10^{-4}	4.43×10^{-4}
		Max-norm	1.86	1.02	1.25	0.97509	2.53×10^{-3}	1.20×10^{-2}
$\eta_d = 1$	Fast P	1-norm	2.00	1.97	1.99	1.00000	7.00×10^{-5}	1.81×10^{-4}
		Max-norm	1.21	1.01	1.12	0.99169	1.98×10^{-3}	5.78×10^{-3}
	S	1-norm	2.01	1.99	1.99	0.99999	2.09×10^{-4}	4.43×10^{-4}
		Max-norm	1.86	1.02	1.25	0.97510	2.53×10^{-3}	1.20×10^{-2}

Chapter 7

RESULTS FOR POROELASTIC-FLUID SYSTEMS ON TWO-DIMENSIONAL MAPPED GRIDS

Having examined the convergence behavior of the two-dimensional mapped grid fluid-poroelastic code for problems with interfaces on rectilinear grids, let us now turn to curvilinear mapped grids. Two types of problem are treated in this chapter. First, a convergence study is performed for a time-harmonic plane wave scattering off an isotropic poroelastic cylinder — a case for which an analytic solution is available — and second, an acoustic pulse is simulated striking a simplified model of a human femur bone.

7.1 *Cylindrical scatterer*

These test cases model acoustic waves scattering off a poroelastic cylinder. The reference solution is for a time-harmonic acoustic plane wave scattering off a cylindrical isotropic poroelastic body; it was obtained by transforming the poroelastic-acoustic system into a set of coupled Helmholtz equations, then expressing their solutions as bi-infinite series of Bessel and Hankel function modes. The full derivation is given in Appendix C. This solution is very similar to that of Laperre and Thys [55], although it is extended here to imperfect hydraulic contact in addition to open or closed pores. It has a variety of interesting properties, including resonance-like behavior at certain frequencies; a future publication is planned exploring its properties in more detail. This work, however, will confine itself to a specific model — a shale cylinder 5 cm in diameter, in a bath of fluid identical to its pore fluid.

In order to choose sensible input frequencies in the face of the complex behavior of the analytical solution, let us first examine the average energy contained in the cylinder over one cycle as a function of frequency, normalized by the average energy contained in the incident wave over the same volume and time. The required integral of the solution in the radial direction is not readily computable analytically, so Romberg quadrature is used instead, applied recursively until the relative difference between the highest-order iterates at successive depths is less than 10^{-9} . The series solution is also truncated to 35 terms in each direction, by which index the series terms have decreased to 10^{-9} or less of their typical magnitudes for low index at all frequencies below the Biot low-frequency validity cutoff, and are decaying rapidly. Figure 7.1 shows this normalized average energy content with and without viscosity, for open, sealed, and imperfect pore conditions. Based on this frequency response plot, three frequencies were selected that showed strong response for each combination of viscosity and interface permeability to test the convergence of the code. These strongly-responding frequencies were chosen because they would be a more rigorous test of the numerical model: the solution is large within the cylinder, coinciding with the most distorted grid cells, and accurate handling of the transfer of energy and momentum across the interface is likely to be more important. Table 7.1 lists the cases chosen, and




Figure 7.1: Frequency response plot for 5 cm diameter cylindrical scatterer composed of shale, with and without viscosity, for several interface discharge efficiencies η_d . The vertical axis is ratio of the average energy contained within the cylinder over one cycle to the average energy contained within an equal volume of the undisturbed incident wave. The circles indicate cases chosen for convergence evaluation of the mapped-grid poroelastic/fluid code. This plot goes up to a maximum frequency of one-third of the cutoff frequency for validity of low-frequency Biot theory; there is no barrier to modeling all the way up to the cutoff frequency, but the response curve becomes increasingly complicated beyond the portion shown here.

Figure 7.2 plots the energy density at the initial time for two selected cases. Similar plots for all cylindrical scatterer cases can be found in Appendix D.

The grid mapping used here deserves some discussion. It is closely related to the square-to-circle mappings of Calhoun, Helzel, and LeVeque [16], but the function defining the radius of curvature of the grid lines has been modified to improve solution quality. Based on experimentation, having concentric grid lines near the surface of the scatterer seems to reduce error, as does having as little grid line curvature as possible and as even a cell size as possible in the interior of the scatterer. Using these considerations as a guide, in the notation of Section 3.2 of [16] the grid mapping used here is defined by $D(d) = r_1 d / \sqrt{2}$, $R(d) = r_1 \left(\frac{9}{10} + d^{19} - \frac{9}{10} d^{20} \right)$. This mapping has $R'(1) = r_1$ and $R''(1) = 0$, giving very nearly concentric grid lines near the scatterer surface, and maintains a grid line radius of

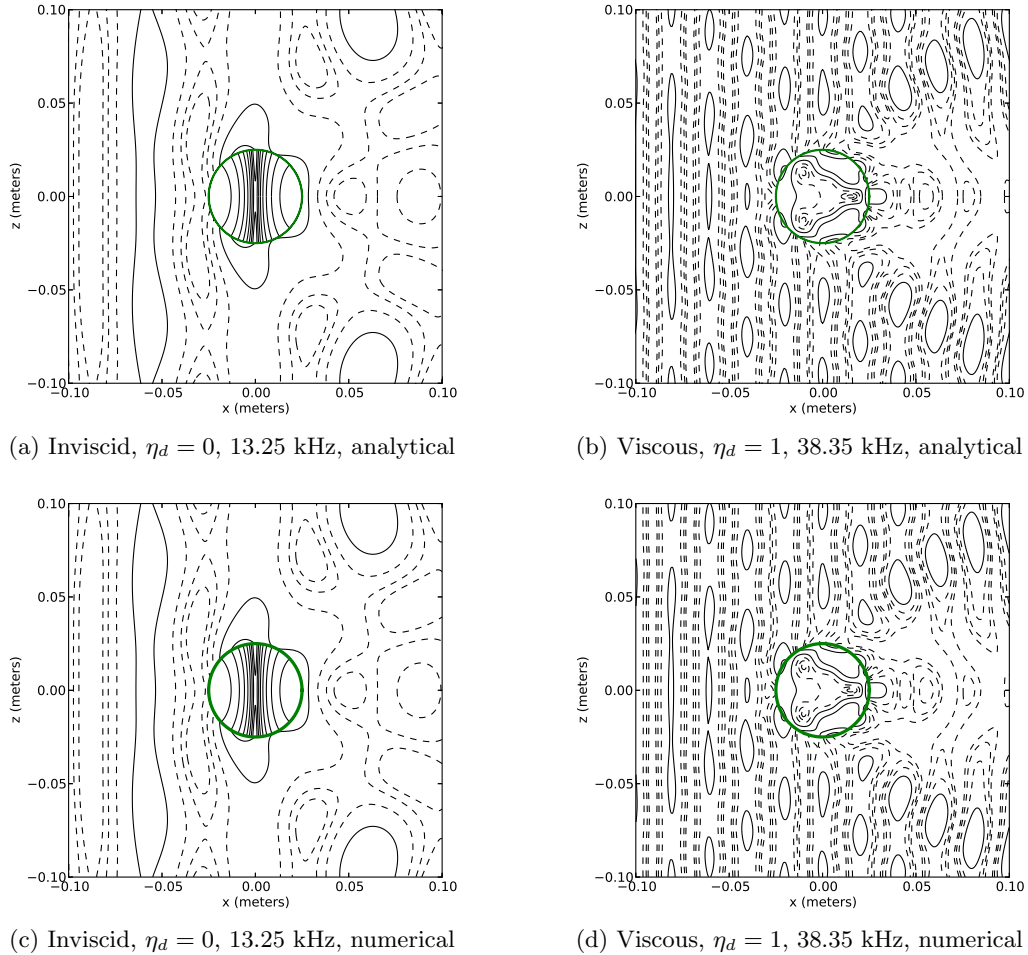




Figure 7.2: Contour plots of energy density for the analytical solutions of selected test cases from Table 7.1 at time $t = 0$, and the numerical solution on a 1024×1024 grid after one cycle. Values are normalized by peak energy density of incident wave. Contours are placed at every power of 2, with dashed contours indicating negative powers; the lowest solid contour is at unit normalized energy density, with the adjacent dashed contour at $\frac{1}{2}$ and the adjacent solid contour at 2. The thick circle is the boundary of the cylinder. Maximum normalized energy densities inside the cylinder in the upper-left plot are 967 at the left edge, and 929 at the right. The numerical solutions closely match the analytical ones. Similar plots for the analytical solutions for all cases can be found in Appendix D.

Table 7.1: Cylindrical scatterer cases chosen for convergence analysis. Note that the reference to “high” frequency here is relative to the scale of Figure 7.1; all cases are well below the cutoff frequency for validity of low-frequency Biot theory.

	Inviscid case frequency (kHz)			Viscous case frequency (kHz)		
	Low	Mid	High	Low	Mid	High
$\eta_d = 0$	13.25	22.25	38.20	15.70	24.55	38.39
$\eta_d = 0.5$	17.25	25.02	39.03	15.80	24.54	38.35
$\eta_d = 1$	17.30	25.09	39.04	15.80	24.54	38.35



(a) Cylindrical scatterer



(b) Femur model

Figure 7.3: Grid mappings used for cylindrical scatterer and femur models, illustrated on a grid 120×120 cells in total size. The thick green circles denote material boundaries. Only the upper-right quadrant of the distorted region is shown; the grid is rectilinear above and to the right of the region shown, while the parts below and to the left are reflections of the upper-right quadrant.

curvature of $\frac{9}{10}r_1$ or greater throughout the interior of the scatterer; the small variation in radius of curvature also helps keep the cell size relatively uniform. Note that even with these modifications, any grid of this type must contain some highly distorted cells, so it should be viewed as one of the worst reasonable cases that might be encountered. Figure 7.3a shows the mapping in action on a coarse grid.

For each case in Table 7.1, the simulation is run for 1.25 cycles of the solution, with initial and boundary conditions set using the true solution as in the previous chapter. The

time step is again chosen so that the global maximum CFL number is 0.9. The error is still measured using the grid energy 1- and max-norms, but since the grid is not uniform an area-weighted grid energy 1-norm is used, computed as

$$\text{error 1-norm} = \frac{\sum_{i=1}^{N_1} \sum_{j=1}^{N_2} \kappa_{ij} \|\mathbf{Q}_{ij,\text{numerical}} - \mathbf{Q}_{ij,\text{true}}\|_E}{\sum_{i=1}^{N_1} \sum_{j=1}^{N_2} \kappa_{ij}}, \quad (7.1)$$

where N_1 and N_2 are the numbers of grid cells in the computational ξ_1 and ξ_2 directions, and κ_{ij} is the cell area ratio discussed in Section 5.1. This is meant to mimic the spatial average of the energy norm of the error. The simulation domain is a square 20 cm on a side, with the cylinder placed at the center. Since the grid mapping is simplest if grid dimensions are used that are a multiple of eight (the ratio of total domain size to cylinder radius), and since the smaller number of cases encourages devoting more computing resources to each one, the grid sizes used to test convergence are increased here to 128×128 , 256×256 , 512×512 , and 1024×1024 .

Figure 7.4 shows the results of these convergence tests. Convergence rates are degraded relative to the rectilinear grid results of Section 6.2. The code achieves roughly first-order convergence at large grid sizes in the 1-norm, with faster convergence on coarser grids; in the max-norm, due to the compounding of error due to the nonsmooth grid mapping with operator splitting error and the omission of the second-order correction term at the surface of the scatterer, convergence rates are typically below first-order, in some cases as low as order $\frac{1}{2}$. Relative errors on the finest grid are typically a few tenths of a percent in the 1-norm, and a few percent in the max-norm. The culprits for these poor convergence results are the highly-distorted, nearly triangular grid cells where the 45° diagonals intersect the surface of the scatterer. These cells have very high error compared to all other cells in the model; moreover, any logically rectangular grid mapping of this type must produce such cells when it maps a rectangle on the computational grid to a circle on the physical grid. A polar grid could be used instead, but it would introduce difficulties at the pole, and would be difficult to generalize to model multiple scatterers or other additional geometry features. In future work, these problems might be ameliorated somewhat by aggressive use of adaptive mesh refinement — even though the convergence rate is slow, the error on these corner cells still decreases as the grid is refined, and heavy grid refinement could still have a reasonable computational cost if it is restricted to the vicinity of a few problem cells. Despite these problems, however, the bottom half of Figure 7.2 shows that on a fine grid, after one cycle of the periodic solution, the numerical solution returns very close to its proper value, without any significant artifacts visible. (Note that the numerical solutions shown in the figure were computed using the MC limiter, as opposed to the convergence results, which were computed without any limiter. The limiter was included because its use would be more typical for solving a problem whose solution is not already known.)

7.2 Femur model

With the results of the previous sections providing reason to be confident in the accuracy of the code, at least at a rough level, let us now turn to a model of a biological system. Specifically, let us examine a highly simplified model of an adult human femur bone. This

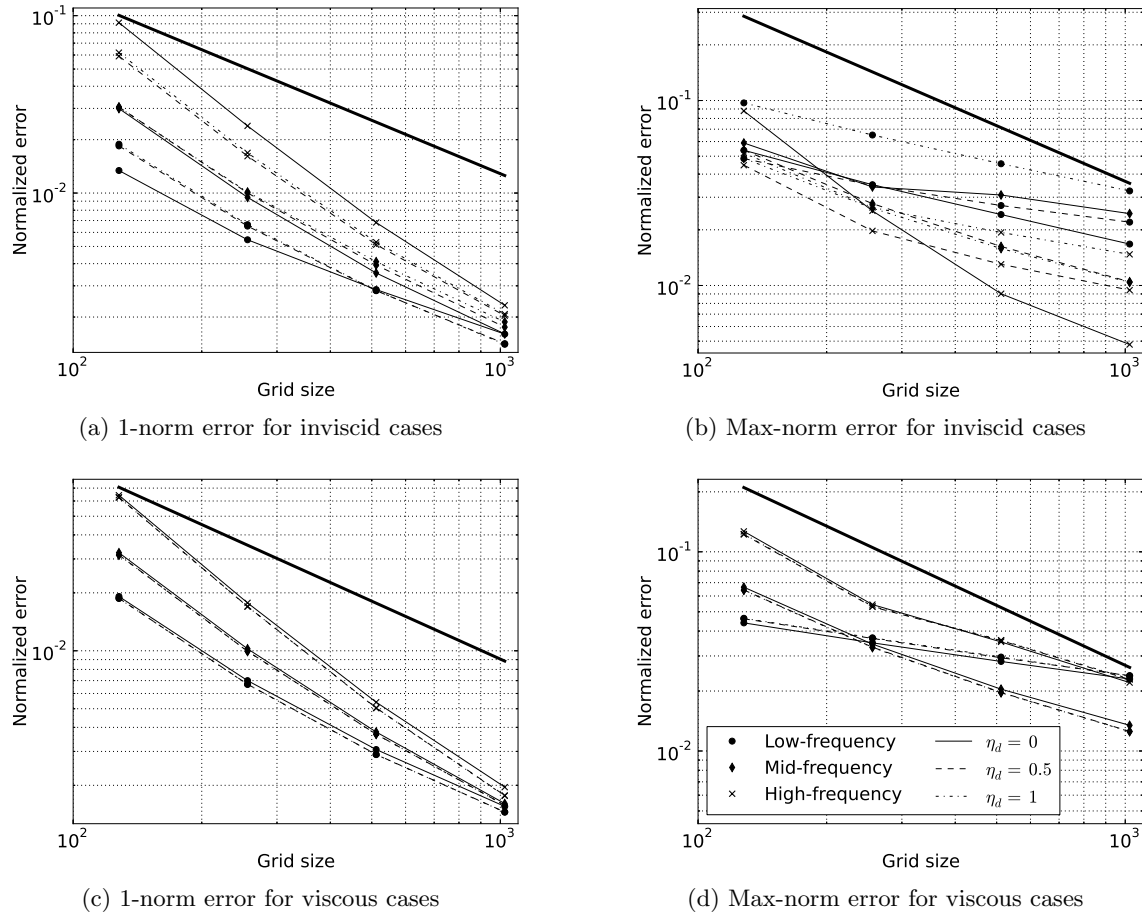


Figure 7.4: Convergence behavior of cylindrical scatterer test cases. The thick black line is a first-order reference line. Error is normalized by the corresponding norm of the true solution on the computational grid.

model is composed of an inner cylinder of marrow-filled cancellous bone of radius 7 mm, surrounded by a concentric cylindrical shell of cortical bone 5 mm thick. The outer shell is in turn surrounded by a fluid bath, with both the inner and outer walls of the shell taken to be impermeable to fluid flow. The properties used for both types of bone material are given in Table A.1; the fluid bath is taken to be water with bulk modulus 2.25 GPa and density 1000 kg/m^3 . The initial condition is an acoustic pulse traveling in the $+x$ direction toward the bone. The pulse has a Gaussian profile with frequency width 100 kHz, and starts with its peak 15 mm away from the surface. While the pulse frequency width is well beyond the low-frequency Biot cutoff frequency, this is intended as a demonstration example only, as opposed to an accurate quantitative model. Since the model is linear, the incoming

pulse is given a peak pressure of 1 Pa for simplicity. This concentric cylindrical geometry is modeled using a modification of the grid mapping used for the cylindrical scatterer, with the grid lines chosen to be concentric circles between the inner and outer radius of the cortical bone shell, and to have a constant radius of curvature inside the cancellous bone core. Figure 7.3b shows the mapping on a coarse grid. The actual computational grid is 800×800 cells. Simple zero-order extrapolation is used to implement nonreflecting boundary conditions; the simulation time is too short for significant interaction between the reflected and transmitted waves from the bone and the boundaries of the domain in any case. The total domain size is 96 mm by 96 mm, with the bone at the center. Since, unlike the other test cases in this chapter, no attempt is being made to isolate the convergence behavior of the wave-propagation algorithm here, the MC limiter was used for all waves in this computation.

Figure 7.5 shows a few relevant quantities from this simulation, at a snapshot 18 μs after the start. The plots show energy density normalized by the peak energy density of the incident pulse, for comparison with the cylindrical scatterer plots, as well as the maximum in-plane shear stress, as a general measure of the deformational load on the bone. Some quantities are also plotted that may be biologically relevant: the fluid pressure in the bone, and the magnitude of fluid flow velocity (the fluid volume flow rate \mathbf{q} divided by porosity). The inclusion of the latter two is motivated by the review paper of Hillsley and Frangos [45], who identify interstitial fluid flow as relevant to osteogenesis. While the plots in Figure 7.5 show some artifacts due to the highly distorted grid mapping — namely, the kinks in some of the contours along the lines of nearly triangular cells extending diagonally from the center of the circles — the solution looks qualitatively good overall. At the point shown in the solution, the leading fast P wavefront has passed entirely through the bone. Meanwhile, separate fast P and S waves from the initial impact of the pulse, visible in Figures 7.5a and 7.5b as the lobes coming off vertically and to the right from the top and bottom of the cancellous bone core, are still propagating through the cortical bone. There are local regions of high fluid velocity (Figure 7.5c), particularly in the cancellous bone but also associated with the fast P wave propagating around the cortical bone; there are also local regions of high or low fluid pressure, associated with a slow P wave to the left of the top and bottom of the cancellous core in Figure 7.5d, and with the reflection of the fast P wave off the right exterior surface of the cortical bone. The response of the pore fluids is hampered, though, by the impermeability of the bone surfaces, which prevents slow P waves from being strongly excited by the incident wave.

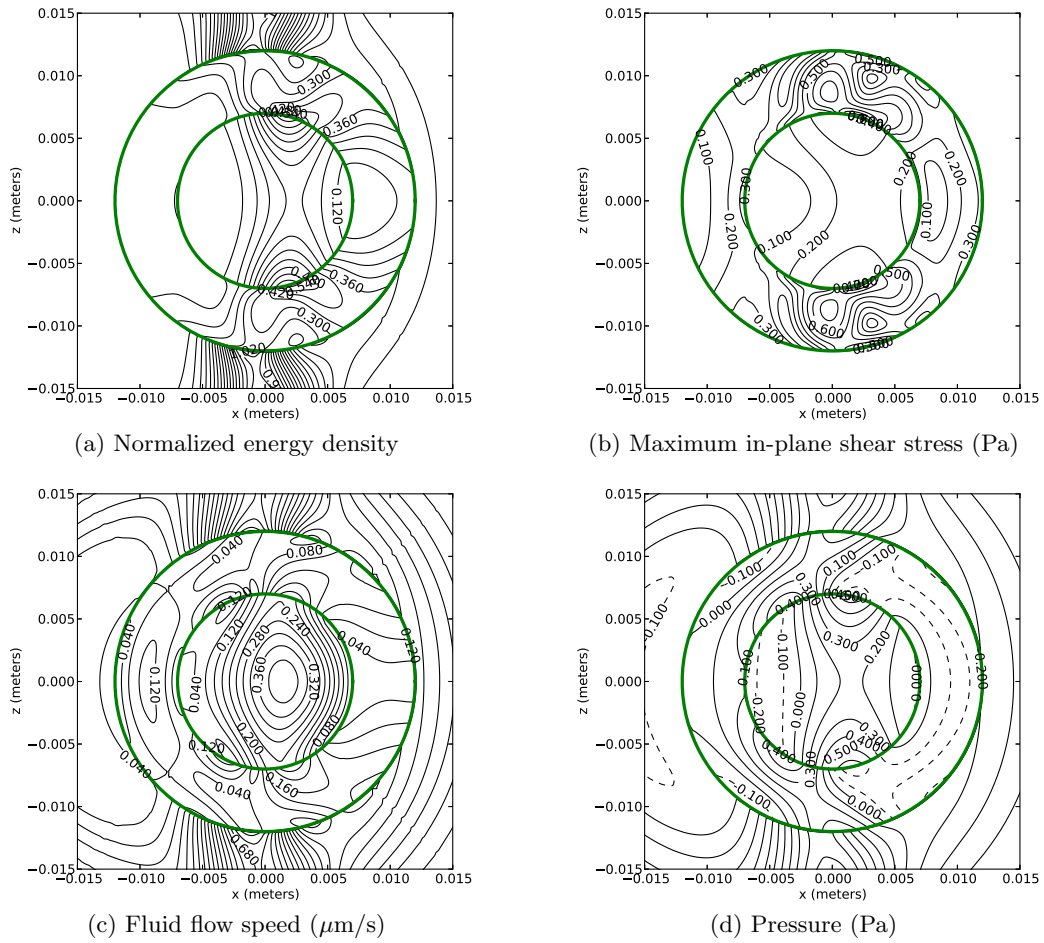


Figure 7.5: Contour plots of various solution quantities in the femur model, $18 \mu\text{s}$ after the start of the simulation. The thick circles mark the boundaries between different media.

Chapter 8

**METHODS FOR POROELASTIC-FLUID SYSTEMS ON
THREE-DIMENSIONAL MAPPED GRIDS**

Despite the increase in dimensionality and the expansion of the state vector from 8 to 13 elements, most of the details of the numerical method for three-dimensional poroelasticity and poroelastic-fluid systems on mapped grids are closely analogous or identical to two dimensions. However, for completeness' sake the changes between the two-dimensional and three dimensional method are described here. Section 8.1 describes how mapped grids are used in three dimensions, while Section 8.2 briefly describes the Euler angles used to relate one set of axes to another. In Section 8.3 the Riemann solution process is discussed, including the interface condition matrices used to solve the Riemann problem between a fluid and a poroelastic medium in three dimensions. In addition, the final steps are also taken to prove that the energy density \mathcal{E} is a strictly convex entropy function for the three-dimensional system, and to exhibit the corresponding reduced system, in Section 8.4.

There are also a few algorithmic details that differ between two dimensions and three. Specifically, the combination of mapped grids and anisotropic materials requires a change to the limiter algorithm in order to cope with potential ambiguities in the eigenvectors associated with the shear waves; this change is discussed and evaluated in two dimensions in Section 8.5. Also, since the transverse Riemann solutions for three-dimensional high-resolution finite volume methods are both computationally expensive and time-consuming to program, dimensional splitting becomes very attractive. While dimensional splitting is by no means a new algorithmic development, Section 8.6 gives a simple overview and compares dimensional splitting with transverse propagation for the viscous, high-frequency plane wave test cases of Chapter 4, in order to evaluate its effects on solution quality. Finally, Section 8.7 gives a brief summary of the software frameworks in which these algorithms are implemented.

8.1 *Mapped grids in three dimensions*

The work of this section was inspired by correspondence with Professor Donna Calhoun, of Boise State University.

As in two dimensions, the three-dimensional numerical solution procedure uses logically rectangular mapped grids. In three dimensions, though, mapped grid quantities such as interface normals and cell face areas are more difficult to define.

Each cell in the mapped grid is defined by the physical coordinates of its vertices, which are computed by applying the mapping function to the vertices of the cell in the computational domain. From these vertices, however, the cell is defined by a trilinear mapping. Let η_1 , η_2 , and η_3 be cell-local computational coordinates, defined from the

global computational coordinates by

$$\eta_i = \frac{\xi_i - \xi_{i0}}{\Delta\xi_i}, \quad (8.1)$$

where ξ_{i0} is the lowest extent of the cell in global computational coordinate i and $\Delta\xi_i$ is the grid spacing in computational coordinate i . In local computational coordinates, the cell is thus the unit cube $[0, 1]^3$. From these coordinates, the cell is parameterized in physical coordinates by

$$\mathbf{r}(\eta_1, \eta_2, \eta_3) = \sum_{i=1}^2 \sum_{j=1}^2 \sum_{k=1}^2 \mathbf{r}_{ijk} N_i(\eta_1) N_j(\eta_2) N_k(\eta_3), \quad (8.2)$$

where the \mathbf{r}_{ijk} are the vertices, with each subscript denoting the position in the corresponding computational direction (so for example \mathbf{r}_{112} is the $-\xi_1, -\xi_2, +\xi_3$ vertex), and the N functions are the standard linear shape functions $N_1(\eta) = 1 - \eta$, $N_2(\eta) = \eta$.

Defining the normal vector to a cell face stretched between four essentially arbitrary vertices is not trivial, because the vertices may not necessarily be coplanar. It seems sensible, though, to require the aggregate normal \mathbf{n} and area A of a face to satisfy the condition

$$\mathbf{n}A = \int_{\text{face}} \mathbf{n}_{\text{local}} dA, \quad (8.3)$$

where $\mathbf{n}_{\text{local}}$ is the local unit normal at each point on the face. If the normals are taken to point outward, this implies that

$$\sum_{i \in \text{faces}} \mathbf{n}_i A_i = \int_{\partial(\text{cell})} \mathbf{n}_{\text{local}} dA = 0, \quad (8.4)$$

a fundamental property of a closed surface. In particular, for a conservative finite volume method, this implies that the net flux of a constant flux vector through the cell is zero, which it should be since the divergence of a constant vector is zero.

The simplest way to satisfy (8.3) is to calculate the integral on the right, then let A be the magnitude of the resulting vector and \mathbf{n} be the unit vector in that direction. For the parameterization (8.2), this integral reduces to the cross product of the vectors connecting the midpoints of the face edges, so, for example, on the $+\xi_1$ face we get

$$\mathbf{n}A = \left(\frac{1}{2} (\mathbf{r}_{221} + \mathbf{r}_{222}) - \frac{1}{2} (\mathbf{r}_{211} + \mathbf{r}_{212}) \right) \times \left(\frac{1}{2} (\mathbf{r}_{212} + \mathbf{r}_{222}) - \frac{1}{2} (\mathbf{r}_{211} + \mathbf{r}_{221}) \right). \quad (8.5)$$

As an aside, the area A calculated by taking the magnitude of this vector is not the full face area; rather, it is the area of the projection of the face onto the plane of its own average normal. This can be seen by taking the dot product of (8.3) with \mathbf{n} :

$$\mathbf{n} \cdot \mathbf{n}A = A = \int_{\text{face}} \mathbf{n} \cdot \mathbf{n}_{\text{local}} dA. \quad (8.6)$$

The cell volume, needed for calculation of the capacity κ , is more difficult to compute. While an analytical expression can be developed for the integral of the Jacobian of the

map (8.2), this expression is quite cumbersome, and it is easier to evaluate the integral by quadrature. The Jacobian is quadratic in each local coordinate η_i , so a tensor product of two-point Gauss-Legendre quadrature in each direction is sufficient to evaluate it exactly to machine precision. The cell centroid location, which is not a fundamental part of the finite volume method but is useful for evaluating spatially-varying initial and boundary conditions, can be calculated the same way — since the position vector (8.2) is first-order in each local coordinate, its product with the Jacobian is at most third-order, so two-point Gauss-Legendre quadrature still suffices to evaluate it exactly.

8.2 Euler angles

The numerical methods here are designed to work with orthotropic poroelastic materials with arbitrarily-oriented principal axes. There are many ways to describe the orientation of one set of axes with respect to another, but here a set of Euler angles are used that are inspired by the commonly-used yaw, pitch, and roll angles of dynamics. Figure 8.1 shows how these angles relate one set of axes to another, and also illustrates the positive sense of the rotations. The new axes x' , y' , and z' are obtained from the old axes x , y , and z by first rotating counterclockwise about the z axis by the yaw angle, then rotating upward (clockwise) about the new y axis by the pitch angle, then finally rotating counterclockwise about new the x axis by the roll angle. Mathematically, the rotation matrix $\boldsymbol{\alpha}$ that transforms a vector from the $x'y'z'$ axes to the xyz axes is

$$\boldsymbol{\alpha} = \mathbf{R}_z(\psi)\mathbf{R}_y(-\phi)\mathbf{R}_x(\theta), \quad (8.7)$$

where ψ is the yaw angle, ϕ is pitch, θ is roll, and $\mathbf{R}_\xi(\delta)$ is the elementary rotation matrix that rotates counterclockwise by an angle δ about the ξ axis.

8.3 Riemann problems on three-dimensional mapped grids

The solution process for Riemann problems on three-dimensional mapped grids is very similar to the process on two-dimensional mapped grids detailed in Section 5.2. This section will primarily focus on the changes necessary in passing to three dimensions.

8.3.1 Eigenvalues and eigenvectors

As in two dimensions, the eigenvectors for three-dimensional acoustics are quite simple; for the matrix $\check{\mathbf{A}} = n_x\mathbf{A} + n_y\mathbf{B} + n_z\mathbf{C}$ they may easily be verified as

$$\begin{aligned} \mathbf{r}_{\text{acoustic, left}} &= (0 \ 0 \ 0 \ 0 \ 0 \ 0 \ -Z_f \ 0 \ 0 \ 0 \ n_x \ n_y \ n_z)^T \quad (\text{left-going wave}) \\ \mathbf{r}_{\text{acoustic, right}} &= (0 \ 0 \ 0 \ 0 \ 0 \ 0 \ Z_f \ 0 \ 0 \ 0 \ n_x \ n_y \ n_z)^T \quad (\text{right-going wave}). \end{aligned} \quad (8.8)$$

For three-dimensional poroelasticity, the eigenvectors may be found by a procedure very similar to that used in two dimensions. Let (n_1, n_2, n_3) be the components of the unit interface normal in the material principal axes. Then, working in principal coordinates, we seek the eigenvalues and eigenvectors of $\check{\mathbf{A}} = n_1\mathbf{A} + n_2\mathbf{B} + n_3\mathbf{C}$. As in two dimensions,




Figure 8.1: Illustration of yaw, pitch, and roll angles. The $x'y'z'$ axes are obtained by rotating the xyz axes through a yaw angle of 60° , a pitch angle of 30° , and a roll angle of 30° . The x axis points out of the page, and all three rotations are in the positive sense.

this eigenproblem can first be converted to a 13×13 symmetric-definite generalized eigenproblem, and then, exploiting the block structure $\check{\mathbf{A}}$, reduced to a 6×6 symmetric-definite eigenproblem,

$$\check{\mathbf{A}}_{sv}^T \mathbf{E}_s \check{\mathbf{A}}_{sv} \mathbf{r}_v = \lambda^2 \mathbf{E}_v \mathbf{r}_v. \quad (8.9)$$

From here, \mathbf{E}_v can again be factorized as $\mathbf{L}\mathbf{L}^T$, where in three dimensions the sparse upper-triangular matrix \mathbf{L} is

$$\mathbf{L} = \begin{pmatrix} \sqrt{\frac{\Delta_1}{m_1}} & 0 & 0 & \frac{\rho_f}{\sqrt{m_1}} & 0 & 0 \\ 0 & \sqrt{\frac{\Delta_2}{m_2}} & 0 & 0 & \frac{\rho_f}{\sqrt{m_2}} & 0 \\ 0 & 0 & \sqrt{\frac{\Delta_3}{m_3}} & 0 & 0 & \frac{\rho_f}{\sqrt{m_3}} \\ 0 & 0 & 0 & \sqrt{m_1} & 0 & 0 \\ 0 & 0 & 0 & 0 & \sqrt{m_2} & 0 \\ 0 & 0 & 0 & 0 & 0 & \sqrt{m_3} \end{pmatrix}. \quad (8.10)$$

As before, this allows the reduction of the generalized eigenproblem (8.9) via the substitution $\mathbf{r}_v = \mathbf{L}^{-T} \mathbf{y}$ to the symmetric ordinary eigenproblem

$$\mathbf{M} \mathbf{y} = \lambda^2 \mathbf{y}, \quad \mathbf{M} := \mathbf{L}^{-1} \check{\mathbf{A}}_{sv}^T \mathbf{E}_s \check{\mathbf{A}}_{sv} \mathbf{L}^{-T}. \quad (8.11)$$

The matrix \mathbf{M} can again be written as a quadratic form in the components of the interface normal, whose matrix coefficients can be pre-computed:

$$\mathbf{M} = n_1^2 \mathbf{M}_{11} + n_2^2 \mathbf{M}_{22} + n_3^2 \mathbf{M}_{33} + n_1 n_2 \mathbf{M}_{12} + n_1 n_3 \mathbf{M}_{13} + n_2 n_3 \mathbf{M}_{23}, \quad (8.12)$$

where

$$\begin{aligned} \mathbf{M}_{11} &= \mathbf{L}^{-1} \mathbf{A}_{sv}^T \mathbf{E}_s \mathbf{A}_{sv} \mathbf{L}^{-T} \\ \mathbf{M}_{22} &= \mathbf{L}^{-1} \mathbf{B}_{sv}^T \mathbf{E}_s \mathbf{B}_{sv} \mathbf{L}^{-T} \\ \mathbf{M}_{33} &= \mathbf{L}^{-1} \mathbf{C}_{sv}^T \mathbf{E}_s \mathbf{C}_{sv} \mathbf{L}^{-T} \\ \mathbf{M}_{12} &= \mathbf{L}^{-1} \mathbf{A}_{sv}^T \mathbf{E}_s \mathbf{B}_{sv} \mathbf{L}^{-T} + \mathbf{L}^{-1} \mathbf{B}_{sv}^T \mathbf{E}_s \mathbf{A}_{sv} \mathbf{L}^{-T} \\ \mathbf{M}_{13} &= \mathbf{L}^{-1} \mathbf{A}_{sv}^T \mathbf{E}_s \mathbf{C}_{sv} \mathbf{L}^{-T} + \mathbf{L}^{-1} \mathbf{C}_{sv}^T \mathbf{E}_s \mathbf{A}_{sv} \mathbf{L}^{-T} \\ \mathbf{M}_{23} &= \mathbf{L}^{-1} \mathbf{B}_{sv}^T \mathbf{E}_s \mathbf{C}_{sv} \mathbf{L}^{-T} + \mathbf{L}^{-1} \mathbf{C}_{sv}^T \mathbf{E}_s \mathbf{B}_{sv} \mathbf{L}^{-T}. \end{aligned} \quad (8.13)$$

At this point the three-dimensional eigensolution process diverges slightly from the two-dimensional one, since in three dimensions \mathbf{A}_{sv} , and thus \mathbf{M} , has two null vectors, not just one. The null space of \mathbf{A}_{sv} corresponds to flow of the pore fluid in the plane parallel to the cell interface, so to find a basis for the null space of \mathbf{M} , first let \mathbf{v}_{01} and \mathbf{v}_{02} be two non-parallel vectors in 3-space that are perpendicular to the interface normal. Now let \mathbf{L}_q be the bottom-right 3×3 block of \mathbf{L} , and define $\mathbf{w}_{0i} = \mathbf{L}_q^T \mathbf{v}_{0i}$, $i = 1, 2$. These \mathbf{w}_0 vectors form a basis for the null space of \mathbf{M} , since

$$\mathbf{M} \begin{pmatrix} 0_{3 \times 1} \\ \mathbf{w}_{0i} \end{pmatrix} = \mathbf{L}^{-1} \check{\mathbf{A}}_{sv}^T \mathbf{E}_s \check{\mathbf{A}}_{sv} \mathbf{L}^{-T} \begin{pmatrix} 0_{3 \times 1} \\ \mathbf{L}_q^T \mathbf{v}_{0i} \end{pmatrix} = \mathbf{L}^{-1} \check{\mathbf{A}}_{sv}^T \mathbf{E}_s \check{\mathbf{A}}_{sv} \begin{pmatrix} 0_{3 \times 1} \\ \mathbf{v}_{0i} \end{pmatrix} = 0. \quad (8.14)$$

Define \mathbf{w}_\perp as a unit vector perpendicular to both \mathbf{w}_{01} and \mathbf{w}_{02} , and from the three \mathbf{w} vectors define the matrices

$$\mathbf{Z} = \begin{pmatrix} \mathbf{I}_{3 \times 3} & 0_{3 \times 1} \\ 0_{3 \times 3} & \mathbf{w}_\perp \end{pmatrix}, \quad \mathbf{Y} = \begin{pmatrix} \mathbf{I}_{3 \times 3} & 0_{3 \times 1} & 0_{3 \times 1} & 0_{3 \times 1} \\ 0_{3 \times 3} & \mathbf{w}_\perp & \mathbf{w}_{01} & \mathbf{w}_{02} \end{pmatrix}. \quad (8.15)$$

These matrices are exactly analogous to the \mathbf{Z} and \mathbf{Y} matrices of Section 5.2.1. In particular, making the variable substitution $\mathbf{y} = \mathbf{Y} (\mathbf{u}^T \quad \mathbf{u}_0^T)^T$, where \mathbf{u} is a 4-vector and \mathbf{u}_0 is a 2-vector, and multiplying (8.11) from the left by \mathbf{Y}^T yields a pair of smaller, decoupled, symmetric eigenproblems,

$$\mathbf{Y}^T \mathbf{M} \mathbf{Y} \begin{pmatrix} \mathbf{u} \\ \mathbf{u}_0 \end{pmatrix} = \begin{pmatrix} \mathbf{Z}^T \mathbf{M} \mathbf{Z} & 0_{4 \times 2} \\ 0_{2 \times 4} & 0_{2 \times 2} \end{pmatrix} \begin{pmatrix} \mathbf{u} \\ \mathbf{u}_0 \end{pmatrix} = \lambda^2 \begin{pmatrix} \mathbf{I}_{4 \times 4} & 0_{4 \times 2} \\ 0_{2 \times 4} & \mathbf{W} \end{pmatrix} \begin{pmatrix} \mathbf{u} \\ \mathbf{u}_0 \end{pmatrix}. \quad (8.16)$$

Here the 2×2 matrix \mathbf{W} comes from the inner products of the \mathbf{w}_0 vectors, $w_{ij} = \mathbf{w}_{0i}^T \mathbf{w}_{0j}$. Since we are only interested in the propagating waves, we can ignore the \mathbf{u}_0 part of this system, and deal only with the 4×4 symmetric eigenproblem

$$\mathbf{M}' \mathbf{u} = \lambda^2 \mathbf{u}, \quad \mathbf{M}' = \mathbf{Z}^T \mathbf{M} \mathbf{Z}. \quad (8.17)$$

Once this eigenproblem has been solved, again using the QR algorithm, the appropriate sign can be chosen for the wave speed λ , and the variable substitutions can be unwound to

recover the wave eigenvector, exactly as in two dimensions. As with the two-dimensional code, the three-dimensional Riemann solver also caches eigensystems for each Riemann problem it solves, and checks whether the next Riemann problem will use the same data; if so, it re-uses the previous eigensystem to speed execution in the common special case of domains that contain large sections with straight grid planes and homogeneous composition.

8.3.2 Solution of the Riemann problem

For the case of identical materials on either side of an interface, \mathbf{E} -orthogonality of the eigenvectors allows easy extraction of the wave strengths from the difference in states, just as in two dimensions. For the case of different materials, with an interface condition between them, the overall solution procedure is the same as in two dimensions, but some of the specifics differ because of the different state vector.

First, note that equations (5.12) through (5.14) are expressed independently of spatial dimension — they are just as valid in three dimensions as in two. Thus the same overall procedure can be used to solve Riemann problems with interface conditions as in Chapter 5, the only differences being the size of the state vector \mathbf{Q} and the number of waves. The only task remaining is to write the matrices \mathbf{C}_l and \mathbf{C}_r corresponding to the interface conditions (2.103) and (2.105).

As in Section 5.2.3, the fluid-poroelastic interface condition will be treated first. Taking the left medium to be poroelastic, a component-by-component accounting of the correspondence between physical variables and the entries of \mathbf{Q} gives

$$\mathbf{C}_{l,\text{poro-fluid}} = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & n_x & n_y & n_z & n_x & n_y & n_z \\ n_x & 0 & 0 & 0 & n_z & n_y & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & n_y & 0 & n_z & 0 & n_x & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & n_z & n_y & n_x & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \eta_d & 0 & 0 & 0 & -Z'n_x & -Z'n_y & -Z'n_z \end{pmatrix} \quad (8.18)$$

$$\mathbf{C}_{r,\text{poro-fluid}} = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & n_x & n_y & n_z \\ 0 & 0 & 0 & 0 & 0 & 0 & -n_x & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & -n_y & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & -n_z & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & \eta_d & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}. \quad (8.19)$$

Here the parameter $Z' := Z_f(1 - \eta_d)$ has been introduced to ease typesetting. The vector $\mathbf{n} = (n_x, n_y, n_z)$ is the unit interface normal pointing from the poroelastic medium into the fluid; Z_f is the fluid acoustic impedance, and η_d is the interface discharge efficiency, defined in Section 2.4. If the poroelastic medium is on the right, as in two dimensions the subscripts l and r may simply be exchanged and the normal \mathbf{n} negated.

For the poroelastic-to-poroelastic interface condition, as in two dimensions, let the quantity $\widehat{\mathbf{q}} \cdot \mathbf{n}$ from (2.103) be expressed as a weighted average of the normal flow rates on both sides of the interface, $\widehat{\mathbf{q}} \cdot \mathbf{n} = (1 - \zeta)\mathbf{q}_l \cdot \mathbf{n} + \zeta\mathbf{q}_r \cdot \mathbf{n}$. Then the interface condition matrices

\mathbf{C}_l and \mathbf{C}_r become

$$\mathbf{C}_{l,\text{poro-poro}} = \begin{pmatrix} n_x & 0 & 0 & 0 & n_z & n_y & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & n_y & 0 & n_z & 0 & n_x & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & n_z & n_y & n_x & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & n_x & n_y & n_z \\ 0 & 0 & 0 & 0 & 0 & 0 & \eta_d & 0 & 0 & 0 & -Z'_l n_x & -Z'_l n_y & -Z'_l n_z \end{pmatrix} \quad (8.20)$$

$$\mathbf{C}_{r,\text{poro-poro}} = \begin{pmatrix} n_x & 0 & 0 & 0 & n_z & n_y & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & n_y & 0 & n_z & 0 & n_x & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & n_z & n_y & n_x & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & n_x & n_y & n_z \\ 0 & 0 & 0 & 0 & 0 & 0 & \eta_d & 0 & 0 & 0 & Z'_r n_x & Z'_r n_y & Z'_r n_z \end{pmatrix}. \quad (8.21)$$

As with the fluid-poroelastic matrices, the parameters $Z'_l := (1 - \zeta)Z_f(1 - \eta_d)$ and $Z'_r := \zeta Z_f(1 - \eta_d)$ have been introduced to ease typesetting. Based on the results of Section 5.2.3, $\zeta = \frac{1}{2}$ is used in all cases without further testing.

Aside from these new interface condition matrices, the Riemann solution procedure in three dimensions is identical to that in two dimensions.

8.4 Entropy function

Following Section 3.3, in three dimensions the matrix $\mathbf{\Pi}$ that maps from the full poroelastic state vector \mathbf{Q} to the vector of conserved quantities of the dissipation part of the system \mathbf{u} is

$$\mathbf{\Pi} := \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & \rho_f/\rho & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & \rho_f/\rho & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & \rho_f/\rho \end{pmatrix}. \quad (8.22)$$

As in two dimensions, the fact that \mathbf{u} is conserved under the action of the dissipation can be seen immediately because $\mathbf{\Pi D} = 0$. The matrix \mathbf{G} that from any \mathbf{u} gives the unique equilibrium \mathbf{Q}_{eq} satisfying both $\mathbf{D Q}_{\text{eq}} = 0$ and $\mathbf{\Pi Q}_{\text{eq}} = \mathbf{u}$ is

$$\mathbf{G} := \begin{pmatrix} \mathbf{I}_{10 \times 10} \\ \mathbf{0}_{3 \times 10} \end{pmatrix}. \quad (8.23)$$

From these two matrices, the reduced system can be formed:

$$\partial_t \mathbf{u} + \mathbf{\Pi} \mathbf{A} \mathbf{G} \partial_x \mathbf{u} + \mathbf{\Pi} \mathbf{B} \mathbf{G} \partial_y \mathbf{u} + \mathbf{\Pi} \mathbf{C} \mathbf{G} \partial_z \mathbf{u} = 0. \quad (8.24)$$

This matrix $\mathbf{\Pi}$ allows us to see that the statements $\mathbf{D}\mathbf{Q} = 0$ and $\mathcal{E}'(\mathbf{Q})^T \mathbf{D}\mathbf{Q} = \mathbf{Q}^T \mathbf{E} \mathbf{D}\mathbf{Q} = 0$ are equivalent to $\mathbf{Q}^T \mathbf{E} = \mathbf{v}^T \mathbf{\Pi}$ in three dimensions for some $\mathbf{v} \in \mathbb{R}^{10}$. As in two dimensions, if $\mathbf{Q}^T \mathbf{E} = \mathbf{v}^T \mathbf{\Pi}$, we immediately have $\mathbf{Q}^T \mathbf{E} \mathbf{D}\mathbf{Q} = \mathbf{v}^T \mathbf{\Pi} \mathbf{D}\mathbf{Q} = 0$ since $\mathbf{\Pi} \mathbf{D} = 0$. Conversely, if $\mathbf{D}\mathbf{Q} = 0$, the form of \mathbf{D} also immediately gives $q_x = q_y = q_z = 0$. Thus in this case the last six components of $\mathbf{Q}^T \mathbf{E}$ are $(\mathbf{Q}^T \mathbf{E})_8 = \rho v_x$, $(\mathbf{Q}^T \mathbf{E})_9 = \rho v_y$, $(\mathbf{Q}^T \mathbf{E})_{10} = \rho v_z$, $(\mathbf{Q}^T \mathbf{E})_{11} = \rho_f v_x$, $(\mathbf{Q}^T \mathbf{E})_{12} = \rho_f v_y$, and $(\mathbf{Q}^T \mathbf{E})_{13} = \rho_f v_z$. This implies that if $\mathbf{D}\mathbf{Q} = 0$, $\mathbf{Q}^T \mathbf{E} = \mathbf{v}^T \mathbf{\Pi}$, with \mathbf{v} given by

$$\mathbf{v}^T = (\mathbf{Q}_s^T \mathbf{E}_s \quad \rho v_x \quad \rho v_y \quad \rho v_z). \quad (8.25)$$

Here the stress parts of \mathbf{Q} and \mathbf{E} have been separated out for convenience.

As in two-dimensional poroelasticity, the above condition implies that the reduced system (8.24) is hyperbolic, and satisfies a nonstrict subcharacteristic condition. Equality can again be realized in the subcharacteristic condition — the example of this found in Section 3.3 carries over to three dimensions — but the argument from two dimensions that this is harmless to the numerical solution also carries over to three dimensions. In addition, the matrix $\mathbf{\Pi} \mathbf{A} \mathbf{G}$ again takes the form of orthotropic elasticity, with the fluid pressure coming along as an additional variable that does not feed back into the elastic variables. The flux Jacobian of the reduced system is given by

$$\mathbf{\Pi} \mathbf{A} \mathbf{G} = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 & -n_1 c_{11}^u & -n_2 c_{12}^u & -n_3 c_{13}^u \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & -n_1 c_{12}^u & -n_2 c_{22}^u & -n_3 c_{23}^u \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & -n_1 c_{13}^u & -n_2 c_{23}^u & -n_3 c_{33}^u \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & -n_3 c_{44} & -n_2 c_{44} \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & -n_3 c_{55} & 0 & -n_1 c_{55} \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & -n_2 c_{66} & -n_1 c_{66} & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & n_1 M \alpha_1 & n_2 M \alpha_2 & n_3 M \alpha_3 \\ -\frac{n_1}{\rho} & 0 & 0 & 0 & -\frac{n_3}{\rho} & -\frac{n_2}{\rho} & 0 & 0 & 0 & 0 \\ 0 & -\frac{n_2}{\rho} & 0 & -\frac{n_3}{\rho} & 0 & -\frac{n_1}{\rho} & 0 & 0 & 0 & 0 \\ 0 & 0 & -\frac{n_3}{\rho} & -\frac{n_2}{\rho} & -\frac{n_1}{\rho} & 0 & 0 & 0 & 0 & 0 \end{pmatrix}. \quad (8.26)$$

8.5 Shear waves and revised limiter algorithm

One subtle difficulty in using high-resolution finite volume methods for three dimensional elasticity and poroelasticity comes with the shear waves. In three dimensions, for any given propagation direction there are two possible polarizations of shear wave, which may or may not have the same speed in an orthotropic material. The possibility of different speeds obliges the finite volume code to treat them as two distinct wave families, but if the speeds are the same, the matrix $\mathbf{\check{A}}$ has a two-dimensional eigenspace, without a natural, intrinsic way to assign waves to one family or the other. This still does not present difficulties in the root Lax-Wendroff method, but applying a limiter to a wave requires comparing it to the

upwind wave in the same family. On a rectilinear grid it would be possible to arbitrarily define (say) vertically-polarized shear waves to be one family and horizontally-polarized ones to be the other, but on a general mapped grid it is impossible to choose polarization directions in a way that smoothly varies over all possible cell interface directions — the popular result that “you can’t comb the hair on a sphere” — so there could be discontinuities in the chosen polarization directions from one Riemann problem to the next. The limiter would see these as solution discontinuities, and would act to suppress the higher-order terms of the method around them, even if the solution were in fact smooth. This combination of possibly anisotropic materials and mapped grids thus makes it challenging to formulate a good wave limiting algorithm.

The solution used here is to exploit the \mathbf{E} -orthogonality of the eigenvectors of $\check{\mathbf{A}}$ to find the component of the upstream waves in the direction of the wave to be limited, rather than using the wave family number. In the classical approach to wave limiting, the wave strength ratio θ for wave m at cell interface $(i - 1/2, j, k)$ is computed as

$$\theta_{\text{classical}} := \frac{(\mathcal{W}_{i-1/2,j,k}^m)^T \mathcal{W}_{I-1/2,j,k}^m}{(\mathcal{W}_{i-1/2,j,k}^m)^T \mathcal{W}_{i-1/2,j,k}^m}, \quad (8.27)$$

where interface $(I - 1/2, j, k)$ is the upwind interface. In its most minimal form, applied only to the shear waves, the new *energy inner product limiter* (\mathbf{E} -limiter for short) replaces the unweighted inner products with energy inner products, and takes the inner product of the shear wave to be limited with the sum of the upwind shear waves; if the shear waves are identified as waves S1 and S2, the wave strength ratio is

$$\theta_{\mathbf{E},\text{shear}} := \frac{(\mathcal{W}_{i-1/2,j,k}^m)^T \mathbf{E} (\mathcal{W}_{I-1/2,j,k}^{\text{S1}} + \mathcal{W}_{I-1/2,j,k}^{\text{S2}})}{(\mathcal{W}_{i-1/2,j,k}^m)^T \mathbf{E} \mathcal{W}_{i-1/2,j,k}^m}. \quad (8.28)$$

For an inhomogeneous domain, \mathbf{E} is the energy density matrix for the medium into which the wave to be limited is propagating. Because the wave eigenvectors are \mathbf{E} -orthogonal, the numerator of the right-hand side of (8.28) gives exactly the component of the upwind shear waves in the direction of the wave being limited, regardless of the choice of eigenvectors and assignment of wave family numbers at the upwind interface. Once the wave strength ratio θ has been calculated, the limiter function $\phi(\theta)$ is applied and the wave is scaled accordingly, exactly as in [57].

There is a potential implementation difficulty with the \mathbf{E} -limiter wave strength as computed in (8.28). As written, the formula for computing θ requires knowing which upstream waves are the shear waves. Normally the shear waves are intermediate in speed between the fast and slow P waves, but that does not necessarily have to be the case. In a medium with an extremely low shear modulus, for example, both shear waves might be slower than the slow P wave. Rather than attempt to determine which wave is which, here the \mathbf{E} -orthogonality is again exploited by adding together all the upstream waves, not just the shear waves. The full form of the \mathbf{E} -limiter wave strength ratio as implemented in the three-dimensional simulation code is then

$$\theta_{\mathbf{E}} := \frac{(\mathcal{W}_{i-1/2,j,k}^m)^T \mathbf{E} \sum_{n,\text{same direction}} \mathcal{W}_{I-1/2,j,k}^n}{(\mathcal{W}_{i-1/2,j,k}^m)^T \mathbf{E} \sum_{n,\text{same direction}} \mathcal{W}_{i-1/2,j,k}^n}, \quad (8.29)$$

Table 8.1: Percentage change in error caused by using the \mathbf{E} -limiter, with ordinary waves or f -waves, on the cylindrical scatterer test cases of Chapter 7. Percentage change is relative to error in solution obtained using the conventional wave strength ratio (8.27). Statistics are computed across all 18 test cases.

	Ordinary waves		f -waves	
	1-norm	Max-norm	1-norm	Max-norm
All grids				
Maximum	+3.42%	+2.86%	+61.95%	+10.24%
Minimum	-8.21%	-21.43%	-2.48%	-22.48%
Mean	-1.97%	-1.54%	+5.77%	-0.28%
Median	-2.05%	-0.62%	+2.40%	-0.22%
Finest grid				
Maximum	+3.21%	+1.60%	+61.95%	+5.41%
Minimum	-7.75%	-8.55%	-2.48%	-9.49%
Mean	-3.05%	-0.63%	+9.55%	-0.35%
Median	-3.65%	+0.23%	+5.63%	+1.23%

where the sums are over waves going in the same direction (left or right) as the wave to be limited. The expression in the denominator is a more efficient way of calculating $(\mathcal{W}_{i-1/2,j,k}^m)^T \mathbf{E} \mathcal{W}_{i-1/2,j,k}^m$ for multiple waves — the two expressions are equal by \mathbf{E} -orthogonality of the waves, but computing $\mathbf{E} \sum \mathcal{W}_{i-1/2,j,k}^m$ (a 13-vector, so still a relatively small amount of data), storing it, and successively computing its inner product with each $\mathcal{W}_{i-1/2,j,k}^m$ requires fewer floating-point operations than computing $(\mathcal{W}_{i-1/2,j,k}^m)^T \mathbf{E} \mathcal{W}_{i-1/2,j,k}^m$ for each wave separately.

The \mathbf{E} -orthogonality of the waves means that the new formula (8.29) for computing the wave strength ratio gives the same results as the original formula (8.27) for fast and slow P waves on grids where successive cell interfaces are parallel. However, because the $\check{\mathbf{A}}$ matrix depends on the normal direction of the grid interface, on more general mapped grids waves in one family may not be \mathbf{E} -orthogonal to waves in different families at the upwind interface. It may in fact be appropriate to include contributions from other wave families in such cases — for instance, a plane wave in a single wave family may have components in multiple wave families if expressed in terms of eigenvectors of an $\check{\mathbf{A}}$ matrix computed for a direction different from its propagation direction. In order to assess the effect of this change in the algorithm, the cylindrical scatterer test cases of Chapter 7 are re-run here with the original limiter formulation and the \mathbf{E} -limiter, both using the MC limiter function for $\phi(\theta)$. All 18 scatterer cases are examined. Because the \mathbf{E} -limiter adds together waves from different families, it seems appropriate to also test it in combination with the f -wave approach of Bale,

LeVeque, Mitran, and Rossmannith [5], to see whether weighting the different waves by their speeds in the sum would give noticeably different results. The original limiter formulation is not run here with f -waves because the poroelastic material in the model is isotropic — the f -wave formulation weights waves by their speeds, and because the wave speeds are the same in all directions, this weighting would have no effect on the wave strength ratios when comparing within the same wave family. Table 8.1 lists the percent changes in error due to incorporating the **E**-limiter with or without f -waves. Adding the **E**-limiter by itself typically gives a modest reduction in error, with some cases seeing substantial reduction and others seeing slight increases. Using the **E**-limiter in combination with f -wave, on the other hand, typically increases the error, sometimes dramatically. Because of this, the 3D simulation code uses the **E**-limiter wave strength ratio (8.29) with ordinary waves, not f -waves.

8.6 Dimensional splitting

Transverse wave propagation in three dimensions is both complex and computationally expensive. Beyond the normal Riemann solve, which is always necessary, the classical 3D transverse propagation approach worked out by Langseth and LeVeque [54] and implemented in CLAWPACK requires eight transverse Riemann solves per cell interface, and in addition eight double-transverse Riemann solves, which provide third-order terms that are necessary for stability. Extending the new two-dimensional transverse solve scheme of Chapter 5 to three dimensions in an analogous fashion would require 16 transverse Riemann solves and perhaps as many as 32 double-transverse solves; including the normal solve, this could be as many as 49 Riemann solves per interface, a prohibitively high computational cost. Most of the computational effort in the poroelastic Riemann solver is in a lengthy setup phase where eigenvectors and coefficient matrices are computed, and since this phase does not depend on the cell states or fluctuations, it would be possible to bundle together all of the transverse or double-transverse solves for a particular interface, and solve them all together for a cost only marginally higher than a single solve. This would reduce the number of times the setup phase is run to three times per cell interface, but it would require a substantial rewrite of the CLAWPACK internals, which would be prohibitively time-consuming and error-prone.

Because of the computational expense of the transverse solves, all three-dimensional results in this work are run using dimensional splitting. For the dimensionally-split approach, rather than solving the normal Riemann problems in all grid directions, computing transverse and double-transverse fluctuations, and subsequently updating the cells to the next time step, instead the normal Riemann problems are solved in only one grid direction. Their solutions are used to update the cells to an intermediate state, and this intermediate state is used to solve the normal Riemann problems in the next grid direction; the results are used to update the cells to a new intermediate state, which is used to solve the normal Riemann problems in the final grid direction and update the cells to the next time step. In combination with Strang splitting for the source term, then, the procedure to advance the solution by Δt from time step n to $n + 1$ with dimensional splitting runs as follows:

1. Advance \mathbf{Q}^n by $\Delta t/2$ under the action of the source term. Call the result \mathbf{Q}^* .

Table 8.2: Summary of convergence results for viscous, high-frequency plane wave test cases in two dimensions using dimensional splitting

Incident wave	Error norm	Convergence rate			Worst R^2 value	Error on 800×800 grid	
		Best	Worst	Mean		Best	Worst
Fast P	1-norm	2.03	1.01	1.19	0.99808	2.61×10^{-5}	1.73×10^{-3}
	2-norm	2.03	1.01	1.19	0.99813	3.16×10^{-5}	1.92×10^{-3}
	Max-norm	2.03	1.00	1.19	0.99608	6.44×10^{-5}	4.17×10^{-3}
S	1-norm	2.01	1.00	1.22	0.99588	2.45×10^{-4}	5.55×10^{-3}
	2-norm	2.01	1.00	1.22	0.99525	2.86×10^{-4}	5.85×10^{-3}
	Max-norm	2.00	0.98	1.24	0.99670	5.29×10^{-4}	9.03×10^{-3}
Slow P	1-norm	2.01	1.00	1.17	0.99987	3.47×10^{-6}	3.65×10^{-2}
	2-norm	2.02	1.00	1.17	0.99982	3.99×10^{-6}	5.55×10^{-2}
	Max-norm	2.00	0.85	1.13	0.99720	6.53×10^{-6}	2.88×10^{-1}

2. Advance \mathbf{Q}^* by Δt using just the Riemann solutions in the grid i direction. Call the result \mathbf{Q}^{**} .
3. Advance \mathbf{Q}^{**} by Δt using just the Riemann solutions in the grid j direction. Call the result \mathbf{Q}^{***} .
4. Advance \mathbf{Q}^{***} by Δt using just the Riemann solutions in the grid k direction. Call the result \mathbf{Q}^{****} .
5. Advance \mathbf{Q}^{****} by $\Delta t/2$ under the source term again. The result is \mathbf{Q}^{n+1} .

While this procedure is formally only first-order accurate, for practical grids and time steps it often displays second-order accuracy.

To assess its effect on solution accuracy for viscous poroelasticity, the viscous high-frequency plane wave test cases of Chapter 4 (Table 4.2) were re-run using dimensional splitting. Table 8.2 displays the results. The dimensionally split solution attains second-order accuracy when the wavevector is aligned with the grid direction, but otherwise degrades to its nominal first-order accuracy, and errors are quite high for the slow P wave. The three-dimensional simulations here nonetheless use dimensional splitting exclusively, because it appears to be the only timely way to obtain numerical solutions for these cases, both in terms of software development time and program execution time.

8.7 Numerical software

The numerical solution techniques described here were implemented using a hybrid of several different versions of the CLAWPACK finite volume software. A pure-Fortran implementation was written to interface with CLAWPACK 4.3, which was the last version before the upcoming CLAWPACK 5.0 to support three-dimensional problems. A hybrid Python-Fortran implementation was also written for PYCLAW [52] in order to be able to use the PETCLAW [2] variant of PYCLAW to run in parallel on large workstation-class computers or clusters.

Chapter 9

**RESULTS FOR POROELASTIC-FLUID SYSTEMS ON
THREE-DIMENSIONAL MAPPED GRIDS**

With the numerical methods formulated for three-dimensional poroelasticity and poroelastic-fluid systems, it is now time to apply these methods to some test problems and see how they perform. While the available time does not permit detailed exploration of many three-dimensional problems, Section 9.1 details the construction of plane wave solutions analogous to those of Section 4.1, and Section 9.2 uses these solutions to examine the convergence behavior of the numerical methods of Chapter 8. The results of a demonstration problem that exercises almost all of the functionality of the three-dimensional code are then presented in section 9.3.

9.1 Analytic plane wave solution

Plane wave solutions for homogeneous orthotropic poroelastic media in three dimensions are generated here in a fashion very similar to the two-dimensional solutions of Section 4.1. The process begins by prescribing a unit vector $\boldsymbol{\ell}$ in the direction of the desired wavevector, an angular frequency ω , the orientation of the principal axes of the material, a desired wave family, and for shear waves a desired polarization direction \mathbf{s} . Given these inputs, first the vectors $\boldsymbol{\ell}$ and \mathbf{s} are transformed into the material principal coordinates to simplify subsequent calculations. Following this, the complex wavevector amplitude k and wave eigenvector \mathbf{v} are obtained by taking an ansatz for the solution \mathbf{Q} of the form

$$\mathbf{Q} = \mathbf{v} \exp(i(k(\ell_1 x_1 + \ell_2 x_2 + \ell_3 x_3) - \omega t)). \quad (9.1)$$

Here ℓ_1 , ℓ_2 , and ℓ_3 are the components of $\boldsymbol{\ell}$ in the material principal coordinates, and x_1 , x_2 , and x_3 are distances along the principal material axes.

Substituting this ansatz into the first-order system for three-dimensional poroelasticity (2.79) results in the eigenproblem

$$-i\omega\mathbf{v} + ik\check{\mathbf{A}}\mathbf{v} = \mathbf{D}\mathbf{v}, \quad (9.2)$$

where as usual $\check{\mathbf{A}} = \ell_1\mathbf{A} + \ell_2\mathbf{B} + \ell_3\mathbf{C}$. Rearranging and rescaling by making the substitution $\mathbf{v} = \mathbf{E}^{-1/2}\mathbf{w}$ and multiplying from the left by $\mathbf{E}^{1/2}$ results in complex symmetric generalized eigenproblem

$$\mathbf{E}^{1/2}\check{\mathbf{A}}\mathbf{E}^{-1/2}\mathbf{w} = k^{-1} \left(\omega\mathbf{I} - i\mathbf{E}^{1/2}\mathbf{D}\mathbf{E}^{-1/2} \right) \mathbf{w}. \quad (9.3)$$

This second form of the eigenproblem is easier to work with numerically — multiplying by the square root and inverse square root of the energy density matrix \mathbf{E} improves the relative

scaling of the components of $\check{\mathbf{A}}$, and letting k^{-1} be the eigenvalue allows the null vectors of $\check{\mathbf{A}}$ to correspond to zero eigenvalues rather than the infinite ones that would result if k were placed in the role of eigenvalue.

From the solutions of the eigenproblem (9.3), the eigenvector \mathbf{w} and complex wavenumber k are extracted that correspond to the desired wave family. The wave eigenvector \mathbf{v} is then computed using $\mathbf{v} = \mathbf{E}^{-1/2}\mathbf{w}$. In the case where a shear wave is requested and the two shear wave speeds are equal, the vector from the two-dimensional shear eigenspace is chosen that has solid velocity as close to parallel to the prescribed polarization direction \mathbf{s} as possible. The vector \mathbf{v} is then normalized to unit \mathbf{E} -norm, and its complex phase is adjusted so that the dot product of its solid velocity component with a reference direction — $\boldsymbol{\ell}$ for fast and slow P waves, \mathbf{s} for shear waves — is pure real and positive. Finally, since the eigenproblem was solved using the system matrices for the principal material axes, \mathbf{v} is transformed back into the global computational axes.

9.2 Plane wave convergence studies

As in Chapter 4 in two dimensions, the three-dimensional code was first tested using these analytical plane wave solutions. Based on the results of Chapter 4, and because of the high computational cost of three-dimensional simulation, only viscous high-frequency test cases are run here. Good convergence behavior for these cases implies that the underlying wave propagation algorithm would also perform well for inviscid cases, and from Section 4.3 we already know to expect first-order convergence for low-frequency viscous cases, regardless of how well the code would perform in other cases. Even with the restriction to viscous high-frequency cases, only a relatively small number of cases are examined in order to keep the computational cost of these convergence studies reasonable.

All test cases are run with the orthotropic, transversely isotropic sandstone of Table A.1, at a frequency of 10 kHz. The computational domain for each case is a cube with its center at the origin, discretized with an equal number of cells in each direction; for most cases the sides of the cube are aligned with the global computational axes, but the grid is rotated for some cases to provide a test of the mapped grid capabilities of the code. For the fast P wave and both S waves, the edge length of the domain is one wavelength of the solution (computed as $2\pi/|\operatorname{Re} k|$ for complex k), and the total simulation time is 1.25 periods of the plane wave. For the slow P wave, the edge length is one decay length of the wave, computed as $1/|\operatorname{Im} k|$, which is substantially less than one wavelength even at this high frequency, and the total simulation time is set to 1.25 times the time for a fast P wave in the material 1-direction to cross the domain. These domain sizes are set substantially smaller than for the analogous two-dimensional cases in order to allow the solution to be resolved on coarser grids, reducing the amount of memory and CPU time required. The simulation time step was again chosen so that the global maximum CFL number was 0.9. For all cases, boundary conditions are implemented as before by filling the ghost cells with the true solution evaluated at cell centroids. Strang splitting is used for the source term, and the hyperbolic part of the system is advanced using dimensional splitting; the overall sequence of steps to advance by one time step is exactly as described in Section 8.6. Limiting

Table 9.1: Summary of plane wave test cases in three dimensions. Within each group, cases are ordered by wave speed, fastest first. Cases are numbered starting at zero for consistency with the Python script that runs them. The components of the ℓ vector are given in grid axes, not in global axes.

Cases	Grid axes			Material axes			ℓ vector
	Yaw	Pitch	Roll	Yaw	Pitch	Roll	
0-3	0°	0°	0°	0°	0°	0°	(1, 0, 0)
4-7	0°	0°	0°	0°	0°	0°	(0, 0, 1)
8-11	30°	20°	10°	0°	0°	0°	(1, 0, 0)
12-15	30°	20°	10°	0°	0°	0°	(0, 1, 0)
16-19	30°	20°	10°	0°	0°	0°	(0, 0, 1)
20-23	0°	0°	0°	30°	20°	10°	(1, 0, 0)
24-27	0°	0°	0°	30°	20°	10°	(0, 1, 0)
28-31	0°	0°	0°	30°	20°	10°	(0, 0, 1)
32-35	0°	0°	0°	0°	0°	0°	$\left(\frac{1}{\sqrt{3}}, \frac{1}{\sqrt{3}}, \frac{1}{\sqrt{3}}\right)$

was not used for any of the tests in this section in order to avoid obscuring the convergence behavior of the underlying wave propagation algorithm.

Table 9.1 lists the plane wave cases by groups of four. Within each group, each wave family is tested, in decreasing order of speed — the first case of each group is the fast P wave, and the last is the slow P wave. Cases 5 and 6 are shear waves in the material’s plane of isotropy, so their polarizations must be explicitly specified; case 5 is polarized with its solid velocity in the x direction, while case 6 is polarized in the y direction. Numbering starts at zero for consistency with the zero-based indexing of the Python script that runs these test cases. Cases 0-7 are the simplest, with neither the grid nor the material principal axes rotated from the global computational axes. Note that these cases only propagate waves in the x and z directions; the y direction would be redundant because the x - y plane is the material’s 1 – 2 plane, in which it is isotropic. Cases 8-19 provide a basic test of the mapped grid capabilities of the simulation code — specifically handling of grid interfaces that are not oriented parallel to the global coordinate planes — while cases 20-31 test correct handling of principal material directions that are not aligned with the global axes. All of the above cases examine only waves propagating in the direction of one of the grid axes, because of the dimensional splitting error results of Section 8.6. Cases 32-35, however, send waves propagating obliquely through the grid, in order to see the full effect of dimensional splitting on accuracy in three dimensions.

Tables 9.2, 9.3, 9.4, and 9.5 list the results of these convergence studies. The convergence behavior of the three-dimensional code is exactly what would be expected from the two-dimensional viscous high-frequency results of Section 4.3 and the dimensional splitting

Table 9.2: Convergence results in the 1-norm for three-dimensional plane wave test cases 0-19 of Table 9.1.

Case	Wave type	Error on grid of size			Convergence rate
		50	100	200	
0	Fast P	1.52×10^{-3}	3.60×10^{-4}	8.82×10^{-5}	2.05
1	S1	3.63×10^{-3}	8.74×10^{-4}	2.16×10^{-4}	2.04
2	S2	4.13×10^{-3}	1.00×10^{-3}	2.46×10^{-4}	2.04
3	Slow P	5.67×10^{-5}	1.40×10^{-5}	3.46×10^{-6}	2.02
4	Fast P	2.92×10^{-3}	7.06×10^{-4}	1.73×10^{-4}	2.04
5	S1	3.68×10^{-3}	8.93×10^{-4}	2.20×10^{-4}	2.03
6	S2	3.68×10^{-3}	8.93×10^{-4}	2.20×10^{-4}	2.03
7	Slow P	1.56×10^{-5}	3.84×10^{-6}	9.50×10^{-7}	2.02
8	Fast P	1.56×10^{-3}	3.80×10^{-4}	9.32×10^{-5}	2.03
9	S1	3.61×10^{-3}	8.75×10^{-4}	2.15×10^{-4}	2.04
10	S2	4.12×10^{-3}	9.97×10^{-4}	2.45×10^{-4}	2.04
11	Slow P	5.48×10^{-5}	1.35×10^{-5}	3.34×10^{-6}	2.02
12	Fast P	1.52×10^{-3}	3.60×10^{-4}	8.81×10^{-5}	2.05
13	S1	3.61×10^{-3}	8.76×10^{-4}	2.15×10^{-4}	2.04
14	S2	4.12×10^{-3}	9.99×10^{-4}	2.45×10^{-4}	2.04
15	Slow P	5.59×10^{-5}	1.38×10^{-5}	3.41×10^{-6}	2.02
16	Fast P	2.86×10^{-3}	6.88×10^{-4}	1.69×10^{-4}	2.04
17	S1	3.74×10^{-3}	9.04×10^{-4}	2.22×10^{-4}	2.04
18	S2	3.69×10^{-3}	8.87×10^{-4}	2.19×10^{-4}	2.04
19	Slow P	2.29×10^{-5}	5.63×10^{-6}	1.40×10^{-6}	2.02

results of Section 8.6: all cases display second-order convergence in both the 1-norm and the max-norm, except for cases 32-35, which involve waves not propagating straight along the grid directions, and display first-order convergence in both norms. Because the solution is always well-resolved, the error values are quite small in all cases; on the 200^3 grid, they only reach as high as 1.5% relative error in the max-norm for the off-axis cases, and 0.058% relative error in the max-norm for the grid-aligned cases.

Table 9.3: Convergence results in the 1-norm for three-dimensional plane wave test cases 20-35 of Table 9.1.

Case	Wave type	Error on grid of size			Convergence rate
		50	100	200	
20	Fast P	1.52×10^{-3}	3.60×10^{-4}	8.81×10^{-5}	2.05
21	S1	3.61×10^{-3}	8.76×10^{-4}	2.15×10^{-4}	2.03
22	S2	4.13×10^{-3}	10.00×10^{-4}	2.46×10^{-4}	2.04
23	Slow P	5.62×10^{-5}	1.38×10^{-5}	3.43×10^{-6}	2.02
24	Fast P	1.52×10^{-3}	3.60×10^{-4}	8.81×10^{-5}	2.05
25	S1	3.61×10^{-3}	8.76×10^{-4}	2.15×10^{-4}	2.03
26	S2	4.13×10^{-3}	10.00×10^{-4}	2.46×10^{-4}	2.04
27	Slow P	5.62×10^{-5}	1.38×10^{-5}	3.43×10^{-6}	2.02
28	Fast P	1.52×10^{-3}	3.60×10^{-4}	8.81×10^{-5}	2.05
29	S1	3.61×10^{-3}	8.76×10^{-4}	2.15×10^{-4}	2.03
30	S2	4.13×10^{-3}	9.99×10^{-4}	2.46×10^{-4}	2.04
31	Slow P	5.62×10^{-5}	1.38×10^{-5}	3.43×10^{-6}	2.02
32	Fast P	1.71×10^{-2}	8.43×10^{-3}	4.19×10^{-3}	1.01
33	S1	2.42×10^{-2}	1.20×10^{-2}	5.95×10^{-3}	1.01
34	S2	2.84×10^{-2}	1.41×10^{-2}	6.99×10^{-3}	1.01
35	Slow P	5.09×10^{-4}	2.53×10^{-4}	1.26×10^{-4}	1.00

Table 9.4: Convergence results in the max-norm for three-dimensional plane wave test cases 0-19 of Table 9.1.

Case	Wave type	Error on grid of size			Convergence rate
		50	100	200	
0	Fast P	3.39×10^{-3}	8.16×10^{-4}	2.09×10^{-4}	2.01
1	S1	8.28×10^{-3}	1.93×10^{-3}	4.83×10^{-4}	2.05
2	S2	9.08×10^{-3}	2.24×10^{-3}	5.29×10^{-4}	2.05
3	Slow P	6.58×10^{-5}	1.68×10^{-5}	4.44×10^{-6}	1.94
4	Fast P	7.05×10^{-3}	1.82×10^{-3}	4.56×10^{-4}	1.97
5	S1	6.97×10^{-3}	1.82×10^{-3}	4.73×10^{-4}	1.94
6	S2	6.99×10^{-3}	1.83×10^{-3}	4.73×10^{-4}	1.94
7	Slow P	2.01×10^{-5}	5.21×10^{-6}	1.57×10^{-6}	1.84
8	Fast P	3.46×10^{-3}	8.85×10^{-4}	2.27×10^{-4}	1.96
9	S1	8.06×10^{-3}	1.99×10^{-3}	4.68×10^{-4}	2.05
10	S2	9.23×10^{-3}	2.33×10^{-3}	5.40×10^{-4}	2.05
11	Slow P	7.61×10^{-5}	2.01×10^{-5}	5.18×10^{-6}	1.94
12	Fast P	3.46×10^{-3}	8.34×10^{-4}	2.13×10^{-4}	2.01
13	S1	8.03×10^{-3}	2.00×10^{-3}	4.67×10^{-4}	2.05
14	S2	9.08×10^{-3}	2.36×10^{-3}	5.73×10^{-4}	1.99
15	Slow P	7.53×10^{-5}	1.97×10^{-5}	5.05×10^{-6}	1.95
16	Fast P	7.03×10^{-3}	1.73×10^{-3}	4.42×10^{-4}	2.00
17	S1	9.52×10^{-3}	2.54×10^{-3}	5.86×10^{-4}	2.01
18	S2	9.84×10^{-3}	2.24×10^{-3}	5.76×10^{-4}	2.05
19	Slow P	3.55×10^{-5}	9.47×10^{-6}	2.81×10^{-6}	1.83

Table 9.5: Convergence results in the max-norm for three-dimensional plane wave test cases 20-35 of Table 9.1.

Case	Wave type	Error on grid of size			Convergence rate
		50	100	200	
20	Fast P	3.44×10^{-3}	8.29×10^{-4}	2.12×10^{-4}	2.01
21	S1	8.18×10^{-3}	2.00×10^{-3}	4.89×10^{-4}	2.03
22	S2	9.21×10^{-3}	2.34×10^{-3}	5.80×10^{-4}	1.99
23	Slow P	7.51×10^{-5}	1.96×10^{-5}	5.00×10^{-6}	1.95
24	Fast P	3.44×10^{-3}	8.29×10^{-4}	2.12×10^{-4}	2.01
25	S1	8.18×10^{-3}	2.00×10^{-3}	4.89×10^{-4}	2.03
26	S2	9.21×10^{-3}	2.34×10^{-3}	5.80×10^{-4}	1.99
27	Slow P	7.51×10^{-5}	1.96×10^{-5}	5.00×10^{-6}	1.95
28	Fast P	3.44×10^{-3}	8.29×10^{-4}	2.12×10^{-4}	2.01
29	S1	8.18×10^{-3}	2.00×10^{-3}	4.89×10^{-4}	2.03
30	S2	9.21×10^{-3}	2.34×10^{-3}	5.80×10^{-4}	1.99
31	Slow P	7.51×10^{-5}	1.96×10^{-5}	5.00×10^{-6}	1.95
32	Fast P	2.30×10^{-2}	1.14×10^{-2}	5.67×10^{-3}	1.01
33	S1	4.41×10^{-2}	2.38×10^{-2}	1.26×10^{-2}	0.91
34	S2	5.44×10^{-2}	2.87×10^{-2}	1.49×10^{-2}	0.93
35	Slow P	2.08×10^{-3}	1.10×10^{-3}	5.87×10^{-4}	0.91

Table 9.6: Domain and mapping parameters for demonstration problem. Note that the domain has horizontal dimensions $L_x/2$ by $L_y/2$.

Parameter	z_0	L_x	L_y	H_x	H_y	ξ_{bot}	ξ_{int}	ξ_{top}	z_{bot}	z_{top}	r_{bot}	r_{top}
Value	0 m	2 m	2 m	$\frac{3L_x}{16\pi}$	$\frac{3L_y}{16\pi}$	0.15	0.6	0.9	-1 m	0.5 m	$\frac{2}{\xi_{\text{bot}}}$	$\frac{2}{1-\xi_{\text{top}}}$

9.3 Demonstration problem

As a final task to demonstrate the numerical methods developed in this thesis, and the capabilities of the code implementing them, this section discusses the simulation of an acoustic pulse in brine striking an undulating bed of orthotropic layered sandstone. The surface of the bed is defined by

$$z_{\text{int}}(x, y) = z_0 + H_x \cos\left(\frac{2\pi x}{L_x}\right) + H_y \cos\left(\frac{2\pi y}{L_y}\right), \quad (9.4)$$

with the parameters z_0 , H_x , L_x , H_y , and L_y as given in Table 9.6. Figure 9.1 shows the surface. Below this z coordinate, the domain is composed of the orthotropic sandstone of Table A.1; above, it is composed of the brine from this sandstone. The curved interface between the two media is incorporated into the model using a mapped grid. At every point within the sandstone, the material's plane of isotropy (the plane of the principal 1-2 axes) is parallel to the tangent plane of the surface above, in order to simulate a bed of stone that has been folded, or deposited on a pre-existing uneven surface. In the simulation code, these variable principal axes are implemented by assigning constant material principal directions to each cell, equal to the directions evaluated at the cell centroid. The interface is taken to have open pores ($\eta_d = 1$ in interface condition (2.105)), and the incoming acoustic pulse propagates straight downward in the $-z$ direction. This problem exercises almost all of the capabilities of the three-dimensional code — it involves mapped grids, an orthotropic material with a variable principal direction, and a fluid-poroelastic interface.

This grid mapping for this problem is defined so that one of computational coordinate surfaces follows the interface, with the rest of the map chosen as a compromise between simplicity, smoothness, and the ability to have a flat grid plane at a useful distance below the interface to easily output slices of the solution for later plotting. The computational domain is the unit cube $[0, 1]^3$; in the xy plane, the problem's symmetry allows the physical domain to be chosen as one quarter of a periodic tile of the surface, $[0, L_x/2] \times [0, L_y/2]$. The grid mapping function in the horizontal axes is a simple scaling, $x := \xi_1 L_x/2$ and $y := \xi_2 L_y/2$, while the mapping function for the z coordinate is defined in terms of the




Figure 9.1: Brine-sandstone interface surface for demonstration problem.

horizontal physical coordinates x and y , and the vertical computational coordinate ξ_3 , as

$$z := \begin{cases} z_{\text{bot}} + \frac{z'_{\text{bot}}}{r_{\text{bot}}} \sinh r_{\text{bot}}(\xi_3 - \xi_{\text{bot}}), & \xi_3 < \xi_{\text{bot}} \\ z_{\text{bot}} + z'_{\text{bot}}(\xi_3 - \xi_{\text{bot}}) + a_{\text{bot}}(x, y) \left(\sqrt{1 + 8 \left(\frac{\xi_3 - \xi_{\text{bot}}}{\xi_{\text{int}} - \xi_{\text{bot}}} \right)^2} - 1 \right), & \xi_{\text{bot}} \leq \xi_3 < \xi_{\text{int}} \\ z_{\text{top}} + z'_{\text{top}}(\xi_3 - \xi_{\text{top}}) + a_{\text{top}}(x, y) \left(\sqrt{1 + 8 \left(\frac{\xi_3 - \xi_{\text{top}}}{\xi_{\text{int}} - \xi_{\text{top}}} \right)^2} - 1 \right), & \xi_{\text{int}} \leq \xi_3 < \xi_{\text{top}} \\ z_{\text{top}} + \frac{z'_{\text{top}}}{r_{\text{top}}} \sinh r_{\text{top}}(\xi_3 - \xi_{\text{top}}), & \xi_3 \geq \xi_{\text{top}}, \end{cases} \quad (9.5)$$

where the derived parameters z'_{bot} , z'_{top} , $a_{\text{bot}}(x, y)$, and $a_{\text{top}}(x, y)$ in the mapping function are defined by

$$z'_{\text{bot}} := \frac{z_0 - H_x - H_y - z_{\text{bot}}}{\xi_{\text{int}} - \xi_{\text{bot}}} \quad (9.6)$$

$$z'_{\text{top}} := \frac{z_{\text{top}} - z_0 - H_x - H_y}{\xi_{\text{top}} - \xi_{\text{int}}} \quad (9.7)$$

$$a_{\text{bot}}(x, y) := \frac{z_{\text{int}}(x, y) - (z_0 - H_x - H_y)}{2} \quad (9.8)$$

$$a_{\text{top}}(x, y) := \frac{z_{\text{int}}(x, y) - (z_0 + H_x + H_y)}{2}. \quad (9.9)$$

The values of the other mapping parameters are given in Table 9.6.

The intent of the mapping (9.5) is to provide uniform grid spacings in the shortest grid columns above and below the interface, in order to prevent cells in these columns from




Figure 9.2: Side views of the mapped grid used for the demonstration problem, displayed on a $30 \times 30 \times 60$ grid. The heavy solid line indicates the sandstone-brine interface; the heavy dashed lines mark the boundaries of the regions where the grid is stretched to move the boundaries outward. Because $H_y = H_x$ and $L_y = L_x$, plots of the xz faces would look identical.

being any smaller than necessary. The parameters z'_{top} and z'_{bot} are rates of change of z with respect to ξ_3 in these shortest columns. Coordinates z_{top} , ξ_{top} , z_{bot} , and ξ_{bot} are the boundaries in physical and computational space of the part of the domain where the solution is considered “interesting.” Beyond these coordinates, the hyperbolic sine term smoothly stretches the grid in the vertical direction to move the boundaries of the computational domain further away from the interface, to improve the performance of the non-reflecting boundary conditions. The stretching will tend to blur the solution in the elongated cells, but this is acceptable because a high-quality solution is not required outside the region between z_{bot} and z_{top} . Note that for $\xi_3 \leq \xi_{\text{bot}}$ and $\xi_3 \geq \xi_{\text{top}}$, surfaces of constant ξ_3 are horizontal planes. This is convenient for outputting a horizontal slice of the solution just below z_{bot} for plotting. In the middle portion of the grid, the mapping is designed to have

a continuous first derivative everywhere except at the sandstone-brine interface, in order to avoid any possible spurious internal reflections that might be caused by a nonsmooth mapping, and to improve accuracy in general. The mapping is allowed to be nonsmooth at the sandstone-brine interface because the second-order correction term is omitted there in any case — accuracy will already be degraded there, and requiring the mapping to be smooth at the interface would result in greater variation in cell size elsewhere. Figure 9.2 shows the resulting grid. Note that the closely-spaced cells above the interface are not problematic for stability — since they are in the brine, not the sandstone, the wave speed within them is the acoustic wave speed of 1550 m/s, whereas the fast P wave speed in the sandstone is always at least 5260 m/s. The smallest vertical dimension of the cells above the interface is still over half the smallest vertical dimension below it, so stability is restricted by the fast P wave in the sandstone, not the acoustic wave in the brine.

By symmetry, the boundary conditions at the lateral faces of the domain are set to be reflective — for the faces parallel to the yz plane, ghost cells are set to the value of the adjacent cell in the computational domain but with τ_{xz} , τ_{xy} , v_x , and q_x negated, while for the faces parallel to the xz plane, τ_{yz} , τ_{xy} , v_y , and q_y are negated. Non-reflecting boundary conditions at the top and bottom face are implemented using zero-order extrapolation, with the elongation of the grid mapping near the top and bottom boundaries used to move these boundaries further away from the interface. Moving the boundaries further away allows the waves generated at the sandstone-brine interface to have angles of incidence closer to normal, which reduces reflections from this simple approach, and also postpones the arrival of waves from within the domain at the boundary. The initial state is set to zero everywhere, except for the incoming plane wave, which is defined by its pressure field,

$$p_{\text{in}}(x, y, z) := \begin{cases} 0.5 \text{ Pa} \left(1 + \cos \left(\frac{2\pi(z - z_{0,\text{wave}})}{\lambda_{\text{wave}}} \right) \right), & |z - z_{0,\text{wave}}| < \lambda_{\text{wave}}/2 \\ 0, & \text{otherwise,} \end{cases} \quad (9.10)$$

where $z_{0,\text{wave}} = z_0 + H_x + H_y + 0.6\lambda_{\text{wave}}$, $\lambda_{\text{wave}} = c_{\text{wave}}/f_{\text{wave}}$, c_{wave} is the sound speed in the brine (1550 m/s), and the fundamental frequency f_{wave} is 10 kHz. To give a downward-propagating acoustic wave, the vertical fluid velocity is set to $q_z = -p_{\text{in}}/Z_f$, where Z_f is the acoustic impedance of the brine. The total simulation time is 400 μs , and the dimensions of the grid used are $300 \times 300 \times 600$ cells. The MC limiter is used with all waves, with the full energy inner product wave strength ratio (8.29).

In order to obtain a solution more quickly, this demonstration problem was run in parallel using PETCLAW. While there is a significant memory overhead associated with the PETSc distributed arrays used by PETCLAW, as of this writing it is the only CLAWPACK variant capable of running in parallel using dimensional splitting, which made it the only practical option in terms of run time for a large three-dimensional problem. With PETCLAW the maximum memory footprint for the $300 \times 300 \times 600$ cell grid was roughly 105 GB, while with Fortran CLAWPACK the memory footprint would have been approximately 26 GB; however, the classic Fortran CLAWPACK code is not yet parallelized as of this writing, and while AMRCLAW is parallelized and would have a similarly low memory footprint, it does not support dimensional splitting and would have required a large amount of additional programmer and CPU time for the transverse solves. For expediency, the problem was run

on an Amazon EC2 CR1 high-memory cluster compute node, with 16 MPI processes; a total of 835 time steps were required, with a run time of 25.5 hours without viscosity included or 26.5 hours with viscosity, giving an aggregate throughput of roughly 30,000 cell time steps per CPU-second on the Intel Xeon E5-2670 CPUs used on this machine.

Figures 9.3-9.11 show the solution at time 399.9 μs , the beginning of the final time step, with Figures 9.3-9.7 showing the solution without viscosity included and Figures 9.8-9.11 showing it with viscosity. The plots show an isometric view of the computational domain, with plots rendered on the xz and yz planes and on the horizontal plane at the z coordinate of the centroids of the first layer of cells below $z = z_{\text{bot}}$; this is the highest layer of cells in the sandstone whose centroids all lie in a horizontal plane, so that a plot rendered on the surface defined by these centroids is easy to interpret. These plots are the most heavily postprocessed solution output in this thesis, so they deserve some explanation. The plotted values on the xz and yz planes are generated by first projecting the values from the centroids of the layer of cells on that side of the domain to the appropriate plane, using the problem symmetry — that is, for values on the xz plane, τ_{yz} , τ_{xy} , v_y , and q_y are set to zero, while for values on the yz plane, τ_{xz} , τ_{xy} , v_x , and q_x are set to zero. The locations associated with these values for plotting purposes are the orthogonal projections of the corresponding cell centroids onto the axis planes. The problem symmetry is also used to extend the computed solution to the lateral corners of the computational domain — values at $(x, y) = (0, 0)$, $(L_x/2, 0)$, and $(0, L_y/2)$ are obtained by copying the values at the nearest cells, then setting all the shear stresses and all the horizontal velocity components to zero. In addition, since computing the energy density requires knowing the material principal directions, the principal directions on the xz and yz faces are computed at the points on these planes associated with the projected values, not the original cell centroids. Values plotted in the horizontal plane near the bottom of the domain are associated with the z coordinate of the cell centroids, which is -1.0014 m; this is the z coordinate of the plane shown.

The solution of this demonstration problem is quite complex, but there are a number of clearly recognizable features. For the inviscid case, Figure 9.3, which shows the energy density, provides a broad view with most of the solution features identifiable. At the bottom, the light arc across the horizontal slice and stretching up into the lower parts of the sides of the domain is the initial fast P wave created when the acoustic wave struck the peak of the sandstone. The additional arc sweeping up and inward from the intersection of the fast P wave with the domain edge is the same initial fast P wave reflected off the boundary. Further inward toward the z axis, the light diagonal bands are shear waves originating from the acoustic wave striking the flanks of the sandstone surface. The shear waves can be identified as such by comparing the energy density in Figure 9.3 with the flow speed in Figure 9.5 — because shear deformation does not change pore volume, features easily visible in the energy density but not in the flow speed may be presumed to be shear phenomena. Just below the surface, the slow P wave is clearly visible as a narrow bright band; because the pore structure is open at the surface, a strong slow P wave is excited by the incident acoustic wave. Finally, above the surface, the incident wave has been reflected and has already partially left the computational domain. No noticeable reflections are present from the nominally non-reflecting boundaries at the top and bottom of the domain.




Figure 9.3: Energy density (J/m³) for the demonstration problem at time 399.9 μs, without viscosity.




Figure 9.4: Fluid pressure (Pa) for the demonstration problem at time $399.9 \mu\text{s}$, without viscosity.




Figure 9.5: Fluid flow speed (the magnitude of volume flow rate vector \mathbf{q} divided by the porosity, in m/s) for the demonstration problem at time $399.9 \mu\text{s}$, without viscosity.




Figure 9.6: Vertical-direction normal stress τ_{zz} (Pa) for the demonstration problem at time $399.9 \mu\text{s}$, without viscosity.




Figure 9.7: Vertical-direction solid velocity v_z (m/s) for the demonstration problem at time $399.9 \mu\text{s}$, without viscosity.




Figure 9.8: Energy density (J/m^3) for the demonstration problem at time $399.9 \mu\text{s}$, with viscosity.



Figure 9.9: Fluid pressure (Pa) for the demonstration problem at time $399.9 \mu\text{s}$, with viscosity.




Figure 9.10: Vertical-direction normal stress τ_{zz} (Pa) for the demonstration problem at time $399.9 \mu\text{s}$, with viscosity.




Figure 9.11: Vertical-direction solid velocity v_z (m/s) for the demonstration problem at time $399.9 \mu\text{s}$, with viscosity.

For the viscous case, the solution is generally similar, but the slow P wave is almost entirely suppressed by the viscous dissipation. It is, however, faintly visible as a band of increased pressure just under the surface in Figure 9.9. Comparing the energy density plots of Figures 9.3 and 9.8, the fast P and S waves are also somewhat dissipated and slowed by viscosity, although the vertical direction stress fields (Figures 9.6 and 9.10) are hardly affected aside from the loss of the slow P wave.

Chapter 10

CONCLUSIONS AND FUTURE WORK

10.1 Summary of results

This thesis has encompassed the development of numerical algorithms, and code implementing those algorithms, for high-resolution finite volume modeling of poroelasticity and poroelastic-fluid systems in two and three dimensions. While the work was instigated as an effort to improve bone modeling for numerical analysis of extracorporeal shock wave therapy, its focus has shifted to more general numerical methods development; however, the methods developed stand ready for application shock wave therapy problems.

The mathematical content of this work began with the development of the governing equations of orthotropic poroelasticity and poroelastic-fluid systems in two and three dimensions in Chapter 2. From the fundamental equations describing the physical processes in a poroelastic medium, the chapter followed Professor M. Yvonne Ou's development of a first-order hyperbolic system of PDEs describing the dynamics of a poroelastic medium in two dimensions, then continued with the author's own work with the energy density functional of the system. This functional was initially developed to address scaling problems encountered in taking the ordinary linear algebraic norms of the chosen state vector, but the associated Hessian matrix \mathbf{E} of the energy density quickly found use in proving the hyperbolicity of the system. The energy density was also revealed as a strictly convex entropy function of the system, in essentially the sense of Chen, Levermore, and Liu [21] (except for a final condition of theirs that was proved later, in chapters 3 and 8), which was a useful theoretical consideration in evaluating the possible effect on the numerical solution of the stiff relaxation term of the poroelasticity system. A convenient result on the equipartition of energy in simple poroelastic waves was also demonstrated, which proved helpful in the construction of Riemann solvers later. Following the development of the two-dimensional theory, an analogous first-order linear hyperbolic system was developed for three-dimensional poroelasticity, and proven to have the same properties as the two-dimensional system. The chapter concluded with a brief discussion of the equations of linear acoustics, which were used to model wave propagation in fluids in this work, and the statement of some interface conditions that might be imposed between a poroelastic medium and a fluid, or between two distinct poroelastic media.

Following this, the numerical work began in Chapter 3 with the formulation of a numerical method for two-dimensional poroelasticity on rectilinear grids. While there was little exceptional about the hyperbolic part of the system — it is larger than typical systems treated with high-resolution finite volume methods, but not qualitatively different — the use of operator splitting for the sometimes-stiff source term deserved special attention. Stiff source terms are a known cause of difficulties in finite volume methods for hyperbolic problems [24, 58], but a conjecture by Pember [66] regarding a *subcharacteristic condition*

that should be satisfied by the reduced system was used as a heuristic indicator that would suggest that qualitatively correct solutions would be obtained if it were satisfied. The final step was taken to prove that the energy density functional was a strictly convex entropy function in the sense of Chen, Levermore, and Liu, following which a result of theirs was used to prove that such a subcharacteristic condition was satisfied. The reduced system of two-dimensional poroelasticity was also exhibited — in the limit of strong dissipation, Biot poroelasticity theory approaches ordinary elastodynamics.

Chapter 4 then went on to investigate the results of using the numerical methods of Chapter 3 on a variety of problems. First, the convergence behavior of the numerical code implementing the methods was investigated for plane waves. Without the viscous dissipation source term, the expected second-order accuracy was observed. With viscosity, however, a stiff regime was identified, wherein convergence rates were degraded to first-order even when Strang splitting (which is nominally second-order accurate) was used. This stiff regime is defined by a solution time step size that is greater than roughly the characteristic decay time associated with the eigenvalues of the dissipation matrix. Outside of the stiff regime, second-order convergence was again obtained. Following the investigation of convergence for plane waves, the code was tested on a set of point source cases in homogeneous media used by Carcione [17] and de la Puente et al. [29], and on a test case used by de la Puente et al. with waves from a point source being reflected off and transmitted through an interface between beds of shale and sandstone. In all cases, the finite volume code reproduced the solutions from the literature well; for the case with the interface between different materials, adaptive mesh refinement was used successfully to capture the wave fronts with reduced computational effort compared to refining the entire domain.

Real-world applications of poroelasticity, however, often do not lend themselves to rectilinear grids. Bones are rarely rectangular, and in geophysical applications, both Earth's surface and the boundaries between rock strata underneath are often not straight. Realistic problems involving poroelasticity also often involve interaction between a poroelastic medium and a fluid, such as ultrasound applied to bones, shock wave therapy, or acoustic waves in the ocean interacting with the sea floor. Because of this, in Chapter 5 the numerical methods of Chapter 3 were extended to logically rectangular mapped grids, which can model moderately complex curved geometries in the same overall algorithmic framework as rectilinear grids, and to systems containing both poroelastic and fluid media, coupled by the interface conditions described in Chapter 2. An eigensolution technique was developed that reduced the task of finding the wave speeds and eigenvectors of the two-dimensional poroelastic system from solving an 8×8 asymmetric eigenproblem to a 3×3 symmetric one, allowing the eigensolver to be built into the Riemann solution process. This was useful in allowing the most accurate Riemann solutions possible for orthotropic poroelastic media on mapped grids. A method was also constructed for solving Riemann problems between domains coupled by nearly arbitrary interface conditions, which was necessary in order to solve at the interface between a poroelastic medium and a fluid, or between two poroelastic media separated by an interface that offers resistance to fluid flow. In addition, the incorporation of fluids into the numerical scheme necessitated modifications to the classical high-resolution finite volume method. A new transverse propagation scheme had to be developed for poroelastic-fluid interfaces in order to avoid contaminating the fluid part of the

domain with state variables that only belonged in the poroelastic part, and the classical second-order correction term had to be dropped at such interfaces for the same reason.

The methods of Chapter 5 were then tested in Chapter 6 on rectilinear-grid problems with interfaces. First a preliminary investigation was conducted of the effects of omitting the second-order correction term along a line representing an interface in a homogeneous domain. The results of this were as good as could be expected: solution accuracy remained at second-order in the 1-norm, with error barely higher than if the second-order term were retained everywhere, and while accuracy was degraded to first-order in the max-norm, the actual error was several times lower than if the second-order term were removed everywhere. Following this, convergence tests were conducted for both acoustic waves in a fluid striking a poroelastic medium, and poroelastic waves reflecting and transmitting across an interface between different poroelastic media; both types of test were conducted for a several different interface permeabilities. Results for all of these reflection/transmission test cases were as expected from the homogeneous domain tests: roughly second-order convergence in the 1-norm, and roughly first-order convergence in the max-norm. It is worth noting that because of the change in material properties at the interface and the resulting kink in the true solution there, these are also the results that would be expected from a classical high-resolution finite volume solution.

With the correctness of the new interface-handling method verified, the methods for handling mapped grids were tested in Chapter 7, for both a cylindrical poroelastic scatterer in a fluid bath (for which an analytical solution is known, and given in Appendix C), and for a simplified model of a human femur bone. The mapped grids used for both of these types of test problems contained highly distorted cells; for the cylindrical scatterer, they degraded the convergence rate to first-order in the 1-norm, and order $\frac{1}{2}$ in the max-norm. However, the magnitude of the relative error remained small — no greater than about 3% in the max-norm, and perhaps 0.25% in the 1-norm, on the finest grid used — and plots of the numerical solution compared well with plots of the analytical solution, with no noticeable numerical artifacts present, so the simulation results remain useful. The femur model showed areas of localized fluid motion and substantial variation of fluid pressure within the bone, both of which may be biologically relevant.

Having successfully developed and tested numerical methods for poroelastic-fluid systems on mapped grids in two dimensions, the size and complexity of the problems treated was greatly increased in Chapter 8 with the extension of these methods to three dimensions. While most of the details of the three-dimensional method were the same as in two dimensions, simply scaled up, additional attention had to be paid to proper formulation of a finite volume method on a three-dimensional mapped grid, and to construction of a limiter algorithm that would correctly limit shear waves in both isotropic and orthotropic media on a mapped grid. Chapter 8 also marked the start of the use of dimensional splitting in this work, and its effect on two-dimensional problems was examined first before it was implemented in three dimensions. Because of the very large size of practical three-dimensional problems, the three-dimensional poroelastic-fluid code implementing the methods of this chapter included integration with the PYCLAW [52] framework, which includes the PET-CLAW [2] package that allows parallelization across supercomputers; while the code has not yet been run on more than a single compute node, doing so should be simple, since in

PETCLAW all of the parallelization tasks are handled by the library framework, rather than the user’s problem-specific code.

Finally, the three-dimensional code developed to implement the methods of Chapter 8 was tested in Chapter 9. Plane wave convergence tests similar to those of Chapter 4 were conducted, though with many fewer test cases due to the long run time and much larger memory footprint of the three-dimensional problems. The results of these convergence tests were exactly as expected — the code displayed second-order accuracy for waves propagating in the grid directions, degrading to first-order for off-axis waves due to dimensional splitting error. The code was also applied to a demonstration problem involving an acoustic pulse in brine striking an undulating bed of porous sandstone; this problem exercised almost all of the code’s functionality, and it gave reasonable results.

10.2 *Opportunities for future work*

There are manifold opportunities for future work based on the research here.

10.2.1 *Improvement of accuracy in the stiff regime*

The stiff regime identified in Section 4.3 is a substantial barrier to accurate simulation when operator splitting is used to model the dissipation term for problems at low frequencies and long length scales. As such, improving accuracy in it is a prime opportunity for future research. Based on the literature for ODEs and hyperbolic PDEs with stiff terms, there seem to be two promising approaches for this:

- Generalization of the Riemann solution process to account for the effect of the dissipation term on the evolution of the waves, building on the work of Hittinger and Roe [46].
- Switching to a semidiscrete rather than a fully-discrete scheme, and using an exponential time-integrator [47] to include the effect of the dissipation term.

Of the two, the exponential integrator approach seems easier to implement — Ketcheson’s SHARPCRAW [53] framework, which is available in PYCLAW, already uses a semidiscrete scheme, so the only thing necessary would be to implement an exponential integrator for SHARPCRAW.

10.2.2 *Extension to higher frequencies*

This work only models low-frequency Biot poroelasticity theory. However, it would be straightforward to extend the code developed here to higher-frequency poroelasticity models by the inclusion of additional memory variables to model the frequency-dependent kernel used for the generalization of Darcy’s law to high frequencies [61]. These memory variables only appear in the source term of the system, so all of the wave propagation machinery developed here could be adapted with minimal or no work; the only portion of the code that would need to change is the source term handling.

10.2.3 *Construction of a generalized second-order correction term for interfaces*

In Section 5.3 it was argued that the standard second-order correction term used in classical high-resolution finite volume methods is inappropriate at an interface between media governed by different PDEs. Lacking any better alternative, the second-order term was simply dropped at such interfaces, resulting in a loss of accuracy. It seems possible to develop a generalization of the classical correction term that would be applicable at interfaces; an approach based on the Immersed Interface Method [59, 72, 73] seems effective, but formulating an appropriate limiting scheme for this case is not trivial.

10.2.4 *Incorporation of nonporous solids*

Addition of nonporous solid media to the poroelastic-fluid solver would allow modeling of geophysical problems with nonporous strata, biomedical problems involving nonporous medical implants, or even buried objects in sea bottom sediment. This addition would require developing appropriate interface conditions analogous to (2.103) and (2.105) that would hold between a nonporous solid and a fluid or a poroelastic medium, and adding calculation of wave speeds and eigenvectors for a nonporous solid to the Riemann solver, neither of which would be particularly challenging.

10.2.5 *Incorporation of nonlinear fluid equations of state*

The original motivating application of this work, extracorporeal shock wave therapy, involves high enough pressures that the nonlinear dependence of fluid pressure on dilatation becomes important. Incorporating this behavior would involve adding capabilities to the Riemann solver to handle this nonlinear behavior for Riemann solves between a pair of fluid cells, as well as between a fluid cell and a poroelastic cell, while continuing to satisfy interface condition (2.105) in the latter case. The nonlinear fluid-fluid solve is a standard task for a high-resolution finite volume code and could be built upon Fagnan's previous work [33]; while the nonlinear fluid-poroelastic solve would be more complex, since only one wave would be nonlinear it would still likely not be extremely difficult.

10.2.6 *Extension of fluctuation-based transverse propagation scheme to three dimensions*

While the fluctuation-based transverse propagation scheme of Section 5.3.1 proved effective in two dimensions, it was not generalized to three dimensions due to both lack of time on the part of the programmer and the expected high computational cost. This generalization, though, would still be of theoretical interest, and it could be of practical use as well, if many transverse solves could be bundled together to reduce the number of times the costly setup phase of the Riemann solver must be executed. This bundling of transverse solves would require substantial rewrites of the `step3` and `flux3` subroutines in CLAWPACK, as well as minor changes to the interface with the poroelastic Riemann solver, but the work could be worthwhile.

BIBLIOGRAPHY

- [1] J. M. Al-Besharah, O. A. Salman, and S. A. Akashah. Viscosity of crude oil blends. *Industrial & engineering chemistry research*, 26(12):2445–2449, 1987.
- [2] A. Alghamdi, A. Ahmadi, D. I. Ketcheson, M. G. Knepley, K. T. Mandli, and L. Dalcin. Petclaw: A scalable parallel nonlinear wave propagation solver for python. In *Proceedings of the 19th High Performance Computing Symposia*, pages 96–103. Society for Computer Simulation International, 2011.
- [3] E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. Du Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen. *LAPACK Users' Guide*. Society for Industrial and Applied Mathematics, Philadelphia, PA, third edition, 1999.
- [4] K. Attenborough, D. L. Berry, and Y. Chen. Acoustic scattering by near-surface inhomogeneities in porous media. Technical report, Defense Technical Information Center OAI-PMH Repository [<http://stinet.dtic.mil/oai/oai>] (United States), 1998.
- [5] D. S. Bale, R. J. LeVeque, S. Mitran, and J. A. Rossmannith. A wave propagation method for conservation laws and balance laws with spatially varying flux functions. *SIAM Journal on Scientific Computing*, 24(3):955–978, 2002.
- [6] M. J. Berger and R. J. LeVeque. Adaptive mesh refinement using wave-propagation algorithms for hyperbolic systems. *SIAM Journal on Numerical Analysis*, 35:2298–2316, 1998.
- [7] M. A. Biot. Theory of propagation of elastic waves in a fluid-saturated porous solid. I. Low-frequency range. *Journal of the Acoustical Society of America*, 28(2):168–178, 1956.
- [8] M. A. Biot. Theory of propagation of elastic waves in a fluid-saturated porous solid. II. Higher frequency range. *Journal of the Acoustical Society of America*, 28(2):179–191, 1956.
- [9] M. A. Biot. Mechanics of deformation and acoustic propagation in porous media. *Journal of Applied Physics*, 33(4):1482–1498, 1962.
- [10] T. Bourbié, O. Coussy, and B. Zinszner. *Acoustics of Porous Media*. Gulf Publishing Company, Houston, Texas, 1987.

- [11] J. L. Buchanan and R. P. Gilbert. Determination of the parameters of cancellous bone using high frequency acoustic measurements. *Mathematical and Computer Modelling*, 45:281–308, 2007.
- [12] J. L. Buchanan and R. P. Gilbert. Determination of the parameters of cancellous bone using high frequency acoustic measurements II: inverse problems. *Journal of Computational Acoustics*, 15(2):199–220, 2007.
- [13] J. L. Buchanan, R. P. Gilbert, and K. Khashanah. Determination of the parameters of cancellous bone using low frequency acoustic measurements. *Journal of Computational Acoustics*, 12(2):99–126, 2004.
- [14] J. L. Buchanan, R. P. Gilbert, A. Wirgin, and Y. S. Xu. *Marine acoustics: direct and inverse problems*. SIAM, Philadelphia, 2004.
- [15] M. Čada and M. Torrilhon. Compact third-order limiter functions for finite volume methods. *Journal of Computational Physics*, 228:4118–4145, 2009.
- [16] D. A. Calhoun, C. Helzel, and R. J. LeVeque. Logically rectangular finite volume grids and methods for circular and spherical domains. *SIAM Review*, 50:723–752, 2008.
- [17] J. M. Carcione. Wave propagation in anisotropic, saturated porous media: plane-wave theory and numerical simulation. *Journal of the Acoustical Society of America*, 99(5):2655–2666, 1996.
- [18] J. M. Carcione. *Wave Fields in Real Media: Wave Propagation in Anisotropic, Anelastic, and Porous Media*. Elsevier, Oxford, 2001.
- [19] J. M. Carcione, C. Morency, and J. E. Santos. Computational poroelasticity – a review. *Geophysics*, 75(5):75A229–75A243, 2010.
- [20] J. M. Carcione and G. Quiroga-Goode. Some aspects of the physics and numerical modeling of Biot compressional waves. *Journal of Computational Acoustics*, 3(4):261–280, 1995.
- [21] G.-Q. Chen, C. D. Levermore, and T.-P. Liu. Hyperbolic conservation laws with stiff relaxation terms and entropy. *Communications in Pure and Applied Mathematics*, 47(6):787–830, 1994.
- [22] G. Chiavassa and B. Lombard. Time domain numerical modeling of wave propagation in 2D acoustic / porous media. *arXiv preprint arXiv:1109.3281*, 2011.
- [23] G. Chiavassa and B. Lombard. Time domain numerical modeling of wave propagation in 2D heterogeneous porous media. *Journal of Computational Physics*, 230(13):5288–5309, 2011.

- [24] P. Colella, A. Majda, and V. Roytburd. Theoretical and numerical structure for reacting shock waves. *SIAM Journal on Scientific and Statistical Computing*, 7(4):1059–1080, 1986.
- [25] S. C. Cowin. Bone poroelasticity. *Journal of Biomechanics*, 32:217–238, 1999.
- [26] S. C. Cowin and L. Cardoso. Fabric dependence of bone ultrasound. *Acta of Bioengineering and Biomechanics*, 12(2), 2010.
- [27] S. C. Cowin and M. M. Mehrabadi. Identification of the elastic symmetry of bone and other materials. *Journal of Biomechanics*, 22(6/7):503–515, 1989.
- [28] N. Dai, A. Vafidis, and E. Kanasevich. Wave propagation in heterogeneous porous media: a velocity-stress, finite-difference method. *Geophysics*, 60:327–340, 1995.
- [29] J. de la Puente, M. Dumbser, M. Käser, and H. Igel. Discontinuous Galerkin methods for wave propagation in poroelastic media. *Geophysics*, 73(5):T77–T97, 2008.
- [30] G. Degrande and G. De Roeck. FFT-based spectral analysis methodology for one-dimensional wave propagation in poroelastic media. *Transport in Porous Media*, 9:85–97, 1992.
- [31] H. Deresiewicz and R. Skalak. On uniqueness in dynamic poroelasticity. *Bulletin of the Seismological Society of America*, 53(4):783–788, 1963.
- [32] E. Detournay and A. H.-D. Cheng. Poroelastic response of a borehole in a non-hydrostatic stress field. *International Journal of Rock Mechanics and Mining Sciences and Geomechanics Abstracts*, 25(3):171–182, 1988.
- [33] K. Fagnan. *High-resolution finite volume methods for extracorporeal shock wave therapy*. PhD thesis, University of Washington, 2010.
- [34] K. Fagnan, R. J. LeVeque, T. J. Matula, and B. MacConaghy. High-resolution finite volume methods for extracorporeal shock wave therapy. In *Hyperbolic Problems: Theory, Numerics, Applications*, pages 503–510. Springer, 2008.
- [35] K. M. Fagnan, R. J. LeVeque, and T. J. Matula. Computational models of material interfaces for the study of extracorporeal shock wave therapy. <http://arxiv.org/abs/1202.5065>, 2011.
- [36] S. Feng and D. L. Johnson. High-frequency acoustic properties of a fluid/porous solid interface. I. New surface mode. *Journal of the Acoustical Society of America*, 74(3):906–914, 1983.

- [37] S. K. Garg, A. H. Nayfeh, and A. J. Good. Compressional waves in fluid-saturated elastic porous media. *Journal of Applied Physics*, 45:1968–1974, 1974.
- [38] R. F. Gibson. *Principles of Composite Material Mechanics*. McGraw-Hill, New York, 1994.
- [39] R. P. Gilbert, P. Guyenne, and M. Y. Ou. A quantitative ultrasound model of the bone with blood as the interstitial fluid. *Mathematical and Computer Modelling*, 55:2029–2039, 2012.
- [40] R. P. Gilbert and Z. Lin. Acoustic field in a shallow, stratified ocean with a poro-elastic seabed. *Zeitschrift für Angewandte Mathematik und Mechanik*, 77(9):677–688, 1997.
- [41] R. P. Gilbert and M. Y. Ou. Acoustic wave propagation in a composite of two different poroelastic materials with a very rough periodic interface: a homogenization approach. *International Journal for Multiscale Computational Engineering*, 1(4):431–440, 2003.
- [42] G. H. Golub and C. F. van Loan. *Matrix Computations*. Johns Hopkins University Press, Baltimore, third edition, 1996.
- [43] B. Gurevich and M. Schoenberg. Interface conditions for Biot’s equations of poro-elasticity. *Journal of the Acoustical Society of America*, 105(5):2585–2589, 1999.
- [44] S. Hassanzadeh. Acoustic modeling in fluid-saturated porous media. *Geophysics*, 56:424–435, 1991.
- [45] M. V. Hillsley and J. A. Frangos. Review: Bone tissue engineering: The role of interstitial fluid flow. *Biotechnology and Bioengineering*, 43(7):573–581, 1994.
- [46] J. A. F. Hittinger and P. L. Roe. Asymptotic analysis of the Riemann problem for constant coefficient hyperbolic systems with relaxation. *Zeitschrift für Angewandte Mathematik und Mechanik*, 84(7):452–471, 2004.
- [47] M. Hochbruck and A. Ostermann. Exponential integrators. *Acta Numerica*, 19(1):209–286, 2010.
- [48] A. Hosokawa and T. Otani. Ultrasonic wave propagation in bovine cancellous bone. *Journal of the Acoustical Society of America*, 101(1):558–562, 1997.
- [49] E. R. Hughes, T. G. Leighton, and P. R. White. Investigation of an anisotropic tortuosity in a biot model ultrasonic propagation in cancellous bone. *Journal of the Acoustical Society of America*, 121(1):568–574, 2007.

- [50] E. Jones, T. Oliphant, P. Peterson, et al. SciPy: Open source scientific tools for Python, 2001–.
- [51] F. Kemm. A comparative study of TVD-limiters – well-known limiters and an introduction of new ones. *International Journal for Numerical Methods in Fluids*, 67:404–440, 2011.
- [52] D. I. Ketcheson, K. Mandli, A. J. Ahmadi, A. Alghamdi, M. Q. de Luna, M. Parsani, M. G. Knepley, and M. Emmett. Pyclaw: Accessible, extensible, scalable tools for wave propagation problems. *SIAM Journal on Scientific Computing*, 34(4):C210–C231, 2012.
- [53] D. I. Ketcheson, M. Parsani, and R. J. LeVeque. High-order wave propagation algorithms for hyperbolic systems. *SIAM Journal on Scientific Computing*, 35(1):A351–A377, 2013.
- [54] J. O. Langseth and R. J. LeVeque. A wave propagation method for three-dimensional hyperbolic conservation laws. *Journal of Computational Physics*, 165:126–166, 2000.
- [55] J. Laperre and W. Thys. Scattering of ultrasonic waves by an immersed porous cylinder. *Acoustics Letters*, 16(1):9–16, 1992.
- [56] G. I. Lemoine, M. Y. Ou, and R. J. LeVeque. High-resolution finite volume modeling of wave propagation in orthotropic poroelastic media. *SIAM Journal on Scientific Computing*, 35(1):B176–B206, 2013.
- [57] R. J. LeVeque. *Finite Volume Methods for Hyperbolic Problems*. Cambridge University Press, New York, 2002.
- [58] R. J. LeVeque and H. C. Yee. A study of numerical methods for hyperbolic conservation laws with stiff source terms. *Journal of Computational Physics*, 86:187–210, 1990.
- [59] Z. Li and K. Ito. *The immersed interface method: Numerical solutions of PDEs involving interfaces and irregular domains*. Society for Industrial and Applied Mathematics, Philadelphia, Pennsylvania, 2006.
- [60] X.-D. Liu and E. Tadmor. Third order nonoscillatory central scheme for hyperbolic conservation laws. *Numerische Mathematik*, 79:397–425, 1998.
- [61] J.-F. Lu and A. Hanyga. Wave field simulation for heterogeneous porous media with singular memory drag force. *Journal of Computational Physics*, 208(2):651–674, 2005.
- [62] B. G. Mikhailenko. Numerical experiment in seismic investigations. *Journal of Geophysics*, 58:101–124, 1985.

- [63] C. Morency and J. Tromp. Spectral-element simulations of wave propagation in porous media. *Geophysical Journal International*, 179:1148–1168, 2008.
- [64] P. M. Morse and H. Feshbach. *Methods of Theoretical Physics*. McGraw-Hill, New York, 1953.
- [65] A. Naumovich. On finite volume discretization of the three-dimensional Biot poroelasticity system in multilayer domains. *Computational Methods in Applied Mathematics*, 6(3):306–325, 2006.
- [66] R. B. Pember. Numerical methods for hyperbolic conservation laws with stiff relaxation I. Spurious solutions. *SIAM Journal on Applied Mathematics*, 53(5):1293–1330, 1993.
- [67] J. H. Rosenbaum. Synthetic microseismograms: logging in porous formations. *Geophysics*, 39(1):14–32, 1974.
- [68] J. E. Santos and E. J. Oreña. Elastic wave propagation in fluid-saturate porous media, part II: The Galerkin procedures. *Mathematical Modeling and Numerical Analysis*, 20:129–139, 1986.
- [69] M. D. Sharma. Wave propagation across the boundary between two dissimilar poroelastic solids. *Journal of Sound and Vibration*, 314:657–671, 2008.
- [70] T. H. Smit, J. M. Huyghe, and S. C. Cowin. Estimation of the poroelastic parameters of cortical bone. *Journal of Biomechanics*, 35:829–835, 2002.
- [71] The CLAWPACK authors. CLAWPACK software. www.clawpack.org.
- [72] C. Zhang. The immersed interface method for elastic wave propagations in heterogeneous materials. Rice University Department of Computational and Applied Mathematics Report TR96-29, 1996.
- [73] C. Zhang and R. J. LeVeque. Immersed interface methods for wave equations with discontinuous coefficients. *Wave Motion*, 25:237–263, 1997.
- [74] C. Zimmerman and M. Stern. Scattering of plane compressional waves by spherical inclusion in a poroelastic medium. *Journal of the Acoustical Society of America*, 94(1):527–536, 1994.

Appendix A

MATERIAL PROPERTIES USED IN THIS WORK

Table A.1 lists base properties of all the poroelastic media used in this work. The properties listed are for orthotropic, transversely isotropic media, and are only those needed for simulation in two dimensions. The additional properties needed for three-dimensional simulation can be derived from the transverse isotropy as

$$c_{22} = c_{11} \tag{A.1}$$

$$c_{23} = c_{13} \tag{A.2}$$

$$c_{44} = c_{55} \tag{A.3}$$

$$c_{66} = \frac{c_{11} - c_{12}}{2} \tag{A.4}$$

$$\kappa_2 = \kappa_1 \tag{A.5}$$

$$T_2 = T_1. \tag{A.6}$$

Table A.2 lists derived properties — wave speeds and characteristic dissipation times — in two dimensions. With transverse isotropy, in three dimensions these properties are the same in principal axis 2 as in axis 1.

Table A.1: Base properties of the poroelastic media used in this work. Stone properties are taken from de la Puente et al. [29]; bone properties are from [25, 48, 49, 70]. Cortical bone properties refer to vascular pore space. The properties of the first sandstone material and the glass/epoxy are orthotropic, but the other materials are isotropic. Subscript numbers indicate principal directions.

	Sandstone (orthotropic)	Glass / epoxy	Sandstone (isotropic)	Shale	Cortical bone	Cancellous bone
K_s (GPa)	80	40	40	7.6	14	18.5
ρ_s (kg/m ³)	2500	1815	2500	2210	1960	1960
c_{11} (GPa)	71.8	39.4	36	11.9	20.6	5.2
c_{12} (GPa)	3.2	1.2	12	3.96	10.6	2.4
c_{13} (GPa)	1.2	1.2	12	3.96	10.6	2.4
c_{33} (GPa)	53.4	13.1	36	11.9	20.6	5.2
c_{55} (GPa)	26.1	3.0	12	3.96	5	1.38
ϕ	0.2	0.2	0.2	0.16	0.04	0.75
κ_1 (10 ⁻¹⁵ m ²)	600	600	600	100	630	7 × 10 ⁶
κ_3 (10 ⁻¹⁵ m ²)	100	100	600	100	630	7 × 10 ⁶
T_1	2	2	2	2	2	1
T_3	3.6	3.6	2	2	2	1
K_f (GPa)	2.5	2.5	2.5	2.5	2.3	2.2
ρ_f (kg/m ³)	1040	1040	1040	1040	1060	990
η (10 ⁻³ kg/m·s)	1	1	1	1	1	40

Table A.2: Derived properties of the materials in Table A.1. Wave speeds are correct in the high-frequency limit. The fast P wave speed is c_{pf} , the S wave speed is c_s , the slow P wave speed is c_{ps} , and the time constant for dissipation is τ_d . Subscript numbers indicate principal directions.

		Sandstone (orthotropic)	Glass / epoxy	Sandstone (isotropic)	Shale	Cortical bone	Cancellous bone
c_{pf1}	(m/s)	6000	5240	4250	2480	3290	3260
c_{pf3}	(m/s)	5260	3580	4250	2480	3290	3260
c_{s1}	(m/s)	3480	1370	2390	1430	1620	1680
c_{s3}	(m/s)	3520	1390	2390	1430	1620	1680
c_{ps1}	(m/s)	1030	975	1020	1130	1123	1480
c_{ps3}	(m/s)	746	604	1020	1130	1123	1480
τ_{d1}	(μ s)	5.95	5.85	5.95	1.25	33	92
τ_{d3}	(μ s)	1.82	1.81	5.95	1.25	33	92

Appendix B

**ANALYTIC SOLUTION PROCEDURE FOR A TIME-HARMONIC
PLANE WAVE TRAIN STRIKING A FLAT INTERFACE**

The procedure to obtain an analytic solution for a sinusoidal plane wave train striking a planar interface between two poroelastic media, or a poroelastic medium and a fluid, starts by first specifying the angular frequency ω of the solution, the unit vector \mathbf{p} in the propagation direction of the incident wave, and the wave family desired if the incident wave is in a poroelastic medium. Calculations are done with a complex-valued solution for convenience; for actual use, only the real part of this solution is kept.

Assume an ansatz for the incident wave of the form

$$\mathbf{Q}_{\text{in}}(\mathbf{x}, t) = \mathbf{V}_{\text{in}} \exp(i(k_{\text{in}}\mathbf{x} \cdot \mathbf{p} - \omega t)), \quad (\text{B.1})$$

Substituting this into the PDE (2.43), which models both poroelasticity and acoustics depending on the choice of coefficient matrices, gives the eigenvalue problem

$$-i\omega\mathbf{V}_{\text{in}} + ik_{\text{in}}(p_x\mathbf{A}_R + p_z\mathbf{B}_R)\mathbf{V}_{\text{in}} = \mathbf{D}_R\mathbf{V}_{\text{in}}. \quad (\text{B.2})$$

Note that here \mathbf{A}_R and \mathbf{B}_R are taken to be in the global x - z axes, rather than the material principal axes. The subscript R denotes the side of the interface corresponding to the incident and reflected waves, as opposed to T , which will denote the side corresponding to the transmitted waves. This eigenproblem is solved for k_{in} and \mathbf{V}_{in} , and the appropriate solution is selected for the desired wave family, if applicable. The eigenvector \mathbf{V}_{in} is typically scaled to have unit energy norm at this point.

With the incident wave solution in hand, the reflected and transmitted waves can be calculated. There should be a reflected or transmitted wave in every wave family in each medium, although some of these may be evanescent waves if the angle of incidence is shallow enough. While evanescent waves decay rapidly away from the interface, they are important near it and must be included in the solution if they occur. Each reflected and transmitted wave is assumed to have a similar complex exponential form to the incoming wave, but with wavevectors whose magnitude and direction are both unknown. Specifically, the ansatz taken for the reflected and transmitted waves is

$$\mathbf{Q}_{(R,T)j}(\mathbf{x}, t) = \mathbf{V}_{(R,T)j} \exp(i(\mathbf{k}_{(R,T)j} \cdot \mathbf{x} - \omega t)), \quad (\text{B.3})$$

where j indexes the wave family and (R, T) may be either R or T — this equation holds for both the reflected and the transmitted waves. The total solution field $\mathbf{Q}(\mathbf{x}, t)$ is then taken to be

$$\mathbf{Q}(\mathbf{x}, t) = \begin{cases} \mathbf{Q}_{\text{in}}(\mathbf{x}, t) + \sum_{j=1}^{N_R} \mathbf{Q}_{Rj}(\mathbf{x}, t), & \mathbf{x} \text{ on incident side of interface,} \\ \sum_{j=1}^{N_T} \mathbf{Q}_{Tj}(\mathbf{x}, t), & \mathbf{x} \text{ on outgoing side of interface,} \end{cases} \quad (\text{B.4})$$

where N_R and N_T are the numbers of reflected and transmitted waves.

At first glance, there seem to be too many unknowns to find a unique solution for each wave, but it is also necessary to satisfy the appropriate interface condition; denoting the incident and outgoing sides of the interface by the subscripts “in” and “out,” this condition becomes $\mathbf{C}_{\text{in}}\mathbf{Q}_{\text{in, total}} = \mathbf{C}_{\text{out}}\mathbf{Q}_{\text{out}}$, where $\mathbf{Q}_{\text{in, total}}$ and \mathbf{Q}_{out} are the limits of the state vector approaching the interface from the incoming and outgoing sides, and \mathbf{C}_{in} and \mathbf{C}_{out} are the interface condition matrices of Section 5.2.3. Writing $\mathbf{Q}_{\text{in, total}}$ and \mathbf{Q}_{out} in terms of individual waves gives

$$\mathbf{C}_{\text{in}} \left(\mathbf{Q}_{\text{in}}(\mathbf{x}, t) + \sum_{j=1}^{N_R} \mathbf{Q}_{Rj}(\mathbf{x}, t) \right) \Big|_{\text{at interface}} = \mathbf{C}_{\text{out}} \sum_{j=1}^{N_T} \mathbf{Q}_{Tj}(\mathbf{x}, t) \Big|_{\text{at interface}}. \quad (\text{B.5})$$

Now, let \mathbf{t} be the unit tangent vector along the interface; without loss of generality assume that the interface runs through the origin of coordinates, and let ξ measure distance along it from the origin. Factoring out the $\exp(-i\omega t)$ time dependence, the interface condition (B.5) becomes

$$\mathbf{C}_{\text{in}} \left(\mathbf{V}_{\text{in}} \exp(ik_{\text{in}}\mathbf{p} \cdot \mathbf{t} \xi) + \sum_{j=1}^{N_R} \mathbf{V}_{Rj} \exp(ik_{Rj} \cdot \mathbf{t} \xi) \right) = \mathbf{C}_{\text{out}} \sum_{j=1}^{N_T} \mathbf{V}_{Tj} \exp(ik_{Tj} \cdot \mathbf{t} \xi). \quad (\text{B.6})$$

For this condition to hold at all points on the interface, it is necessary to have $\mathbf{k}_{(R,T)j} \cdot \mathbf{t} = k_{\text{in}}\mathbf{p} \cdot \mathbf{t}$ for every reflected and transmitted wave; denote this common tangential component of the wavevector by k_t , and note that it is entirely specified by the direction and frequency of the incident wave.

Expressing (2.43) in coordinates tangential and normal to the interface, then applying it to the ansatz (B.3) and dividing out the exponential time and space dependence, we obtain for each outgoing wave

$$-i\omega\mathbf{V}_{(R,T)j} + ik_t\check{\mathbf{A}}_{(R,T)}\mathbf{V}_{(R,T)j} + ik_{(R,T)n_j}\check{\mathbf{B}}_{(R,T)}\mathbf{V}_{(R,T)j} = \mathbf{D}_{(R,T)}\mathbf{V}_{(R,T)j}. \quad (\text{B.7})$$

Here $\check{\mathbf{A}}_{(R,T)} = t_x\mathbf{A}_{(R,T)} + t_z\mathbf{B}_{(R,T)}$ is the flux Jacobian in the tangential direction, and $\check{\mathbf{B}}_{(R,T)} = -t_z\mathbf{A}_{(R,T)} + t_x\mathbf{B}_{(R,T)}$ is the flux Jacobian in the direction normal to the interface; $k_{(R,T)n_j}$ is the component of the wavevector in the interface normal direction. The subscript (R, T) is included as a reminder that the system matrices will in general be different on either side of the interface. Multiplying from the left by $\mathbf{E}_{(R,T)}$ to symmetrize and reorganizing, a complex symmetric generalized eigenproblem can be obtained for each $k_{(R,T)n_j}$ and its corresponding $\mathbf{V}_{(R,T)j}$:

$$\left(\omega\mathbf{E}_{(R,T)} - k_t\mathbf{E}_{(R,T)}\check{\mathbf{A}}_{(R,T)} - i\mathbf{E}_{(R,T)}\mathbf{D}_{(R,T)} \right) \mathbf{V}_{(R,T)j} = k_{(R,T)n_j}\mathbf{E}_{(R,T)}\check{\mathbf{B}}_{(R,T)}\mathbf{V}_{(R,T)j}. \quad (\text{B.8})$$

Because these represent reflected and transmitted waves, choose the eigenvalues $k_{(R,T)n_j}$ are chosen that correspond to waves propagating away from the interface (for pure real eigenvalues, for which propagation direction is determined from the sign of the normal energy flux $\mathbf{Q}^H\mathbf{E}\check{\mathbf{B}}\mathbf{Q}$), or that decay away from the interface (for complex eigenvalues). As

a practical note, $\check{\mathbf{B}}$ is always singular, and for the formulation of acoustics used here the matrix on the left-hand side is also singular with some of its null space in common with $\check{\mathbf{B}}$, so this eigenproblem presents numerical difficulties (see, for example, Section 7.7.3 of Golub and van Loan [42]). To obtain an accurate solution, it was necessary to make an additional change of variables to eliminate the null spaces of both matrices, then solve a smaller generalized eigenproblem with both matrices nonsingular. For the actual eigensolution, the `eig` command in Scipy [50] was used, which in turn calls the LAPACK [3] routine `ZGGEV`.

Combined with k_t , these eigensolutions fully define the reflected and transmitted waves of (B.3) up to a scalar factor; because that factor has not yet been determined, the eigenvectors found from (B.8) will be denoted by $\mathbf{v}_{(R,T)j}$, and actual reflected and transmitted wave eigenvectors will be $\mathbf{V}_{(R,T)j} = \beta_{(R,T)j} \mathbf{v}_{(R,T)j}$. To find these scalar factors, return to (B.6) and divide out the common factor of $\exp(ik_t \xi)$ to obtain

$$\mathbf{C}_{\text{in}} \left(\mathbf{V}_{\text{in}} + \sum_{j=1}^{N_R} \beta_{Rj} \mathbf{v}_{Rj} \right) = \mathbf{C}_{\text{out}} \sum_{j=1}^{N_T} \beta_{Tj} \mathbf{v}_{Tj}. \quad (\text{B.9})$$

Rearranging and casting in matrix form results in a linear system for the β values,

$$\left(-\mathbf{C}_{\text{in}} \mathbf{v}_{R1} \quad \dots \quad -\mathbf{C}_{\text{in}} \mathbf{v}_{RN_R} \quad \mathbf{C}_{\text{out}} \mathbf{v}_{T1} \quad \dots \quad \mathbf{C}_{\text{out}} \mathbf{v}_{TN_T} \right) \begin{pmatrix} \beta_{R1} \\ \vdots \\ \beta_{RN_R} \\ \beta_{T1} \\ \vdots \\ \beta_{TN_T} \end{pmatrix} = \mathbf{C}_{\text{in}} \mathbf{V}_{\text{in}}. \quad (\text{B.10})$$

Solving this system completes the information necessary to describe the reflected and transmitted wave fields, and allows computation of the full solution field at any point and time.

Appendix C

**TIME-HARMONIC ANALYTIC SOLUTION FOR AN ACOUSTIC
PLANE WAVE SCATTERING OFF AN ISOTROPIC POROELASTIC
CYLINDER**

This appendix gives the derivation of an analytic solution for a time-harmonic acoustic plane wave impinging on a poroelastic cylinder (or disc in two dimensions). The disc is assumed to be composed of an *isotropic* poroelastic medium. Essentially all of the work in sections C.1 through C.4, along with sections C.5.1, C.5.3, and the ideas of Section C.5.4, are from Professor M. Yvonne Ou of the University of Delaware, and are simply transcribed here. The remainder of this appendix is the author's original work.

The following references are useful:

- Zimmerman and Stern's 1993 scattering paper [74]
- Biot's 1962 poroelasticity paper [9]
- The book *Methods of Theoretical Physics* by Morse and Feshbach [64]

C.1 Preliminaries

First, given the assumption of isotropy, the poroelastic constants can be simplified to

$$\begin{aligned}
 c_{12} &= c_{13} = \lambda \\
 c_{11} &= c_{33} = \lambda + 2\mu \\
 c_{55}^u &= c_{55} = \mu \\
 c_{11}^u &= c_{33}^u = H = K_m + \frac{4}{3}\mu + \alpha^2 M \\
 T_1 &= T_3 = T \\
 m_1 &= m_3 = m.
 \end{aligned} \tag{C.1}$$

Using the notation of [9], \mathbf{U} is the pore fluid displacement and \mathbf{u} is the solid displacement. The vector \mathbf{q} is related to these by

$$\mathbf{q} = \phi \left(\dot{\mathbf{U}} - \dot{\mathbf{u}} \right), \tag{C.2}$$

or

$$\begin{aligned}
 \mathbf{w} &= \phi (\mathbf{U} - \mathbf{u}) \\
 \mathbf{q} &= \dot{\mathbf{w}}.
 \end{aligned} \tag{C.3}$$

The second of Biot's equations, in (8.1) of [9], comes from the following procedures:

1. Rewrite (2.12) as

$$p = -\alpha_1 M \partial_x u_x - \alpha_3 M \partial_z u_z - M \operatorname{div} \mathbf{w}. \quad (\text{C.4})$$

2. Substitute p into (2.11) in our notes to obtain the second equation in [9],

$$\nabla(Qe + R\epsilon) = \rho_{12}\ddot{\mathbf{u}} + \rho_{22}\ddot{\mathbf{U}} - b(\dot{\mathbf{u}} - \dot{\mathbf{U}}). \quad (\text{C.5})$$

In terms of the variables of Section 2.1 and the displacements \mathbf{u} and \mathbf{U} ,

$$\begin{aligned} e &= \operatorname{div} \mathbf{u} \\ \epsilon &= \operatorname{div} \mathbf{U} \\ Q &= \phi M(\alpha - \phi) \\ R &= \phi^2 M \\ \rho_{11} &= (1 - \phi)\rho_s + m\phi^2 - \phi\rho_f \\ \rho_{12} &= \phi\rho_f(1 - T) \\ \rho_{22} &= m\phi^2 \\ b &= \frac{\phi^2 \eta}{\kappa}. \end{aligned} \quad (\text{C.6})$$

C.2 Derivation of a second-order system

To obtain the first equation of the Biot system (equation (8.1) of [9]), write the total stress τ_{ij} as

$$\tau_{ij} = \sigma_{ij} + \delta_{ij}\sigma, \quad \sigma = Qe + R\epsilon. \quad (\text{C.7})$$

Here σ is related to the pore pressure p_f by $\sigma = -\phi p_f$. Substituting this into the equations of motion (2.35) through (2.38) gives

$$\partial_j \tau_{ij} = \partial_j(\sigma_{ij} + \delta_{ij}\sigma) = \rho \partial_t v_i + \rho_f \partial_t q_i = \rho \ddot{u}_i + \rho_f \phi(\ddot{U}_i - \ddot{u}_i), \quad (\text{C.8})$$

or

$$\partial_j \sigma_{ij} = \rho \ddot{u}_i + \rho_f \phi(\ddot{U}_i - \ddot{u}_i) - \partial_i \sigma. \quad (\text{C.9})$$

But from equation (C.5) of the previous section, we know that

$$\partial_i \sigma = \rho_{12} \ddot{u}_i + \rho_{22} \ddot{U}_i - b(\dot{u}_i - \dot{U}_i), \quad (\text{C.10})$$

so

$$\begin{aligned} \partial_j \sigma_{ij} &= \rho \ddot{u}_i + \rho_f \phi(\ddot{U}_i - \ddot{u}_i) - \rho_{12} \ddot{u}_i - \rho_{22} \ddot{U}_i + b(\dot{u}_i - \dot{U}_i) \\ &= (\rho - \rho_f \phi - \rho_{12}) \ddot{u}_i + b(\dot{u}_i - \dot{U}_i) + (\rho_f \phi - \rho_{22}) \ddot{U}_i. \end{aligned} \quad (\text{C.11})$$

Note that the bulk density of the medium is $\rho = (1 - \phi)\rho_s + \phi\rho_f$. It can easily be verified that $\rho - \rho_f \phi - \rho_{12} = \rho_{11}$ and $\rho_f \phi - \rho_{22} = \rho_{12}$, so (C.11) becomes

$$\partial_j \sigma_{ij} = \rho_{11} \ddot{u}_i + \rho_{12} \ddot{U}_i + b(\dot{u}_i - \dot{U}_i). \quad (\text{C.12})$$

To obtain the first Biot equation, note that $\sigma_{ij} = \tau_{ij} - \delta_{ij}\sigma$. From equation (3.28) of [9], we have

$$\sigma_{ij} = 2Ne_{ij} + \delta_{ij}(Ae + Q\epsilon), \quad (\text{C.13})$$

where the parameters N and A are related to the material parameters used thus far by

$$\begin{aligned} N &= \mu \\ A &= \lambda + M(\alpha - \phi)^2. \end{aligned} \quad (\text{C.14})$$

Thus we get

$$\begin{aligned} \partial_j \sigma_{ij} &= \partial_j (2Ne_{ij} + \delta_{ij}(Ae + Q\epsilon)) \\ &= \partial_j (2\mu e_{ij}) + \partial_i (Ae + Q\epsilon) \\ &= \partial_j \left(2\mu \frac{1}{2} (\partial_j u_i + \partial_i u_j) \right) + \partial_i (Ae + Q\epsilon) \\ &= \mu (\partial_j \partial_j u_i + \partial_i \partial_j u_j) + \partial_i (Ae + Q\epsilon). \end{aligned} \quad (\text{C.15})$$

Noting that $\partial_j \partial_j u_i = \Delta u_i$, the Laplacian of component i of the solid displacement, and $\partial_i \partial_j u_j = \partial_i \operatorname{div} \mathbf{u}$, we obtain Biot's first equation,

$$\mu \Delta \mathbf{u} + \nabla ((\mu + A)e + Q\epsilon) = \rho_{11} \ddot{\mathbf{u}} + \rho_{12} \ddot{\mathbf{U}} + b(\dot{\mathbf{u}} - \dot{\mathbf{U}}). \quad (\text{C.16})$$

C.3 Further reduction of unknowns

Similarly to [74], the number of unknowns can be further reduced from 4 to 3. This process starts from the second-order system,

$$\begin{aligned} \mu \Delta \mathbf{u} + \nabla ((\mu + A)e + Q\epsilon) &= \rho_{11} \ddot{\mathbf{u}} + \rho_{12} \ddot{\mathbf{U}} + b(\dot{\mathbf{u}} - \dot{\mathbf{U}}) \\ \nabla (Qe + R\epsilon) &= \rho_{12} \ddot{\mathbf{u}} + \rho_{22} \ddot{\mathbf{U}} - b(\dot{\mathbf{u}} - \dot{\mathbf{U}}). \end{aligned} \quad (\text{C.17})$$

(Remember here that $e = \operatorname{div} \mathbf{u}$ and $\epsilon = \operatorname{div} \mathbf{U}$.) From here, the fluid displacement \mathbf{U} will be replaced with the pore pressure p_f , referred to as p above.

Fluid pressure is related to the solid and fluid dilatations e and ϵ by

$$p = -\frac{1}{\phi}(Qe + R\epsilon) = -\frac{\sigma}{\phi}. \quad (\text{C.18})$$

For the time-harmonic case, make the identification

$$\mathbf{u}(\mathbf{x}, t) = \mathbf{u}(\mathbf{x})e^{-i\omega t}, \quad \mathbf{U}(\mathbf{x}, t) = \mathbf{v}(\mathbf{x})e^{-i\omega t}. \quad (\text{C.19})$$

The second-order system (C.17) then becomes

$$\mu \Delta \mathbf{u} + \nabla ((\mu + A)e + Q\epsilon) + \omega^2 (\rho'_{11} \mathbf{u} + \rho'_{12} \mathbf{v}) = 0 \quad (\text{C.20})$$

$$\nabla (Qe + R\epsilon) + \omega^2 (\rho'_{12} \mathbf{u} + \rho'_{22} \mathbf{v}) = 0, \quad (\text{C.21})$$

where the new complex-valued density parameters ρ'_{11} , ρ'_{12} , and ρ'_{22} are given by

$$\rho'_{11} = \rho_{11} + \frac{ib}{\omega}, \quad \rho'_{12} = \rho_{12} - \frac{ib}{\omega}, \quad \rho'_{22} = \rho_{22} + \frac{ib}{\omega}. \quad (\text{C.22})$$

Since $\sigma = Qe + R\epsilon$, (C.21) can be rewritten as

$$\nabla\sigma + \omega^2 (\rho'_{12}\mathbf{u} + \rho'_{22}\mathbf{v}) = 0, \quad (\text{C.23})$$

which gives

$$\mathbf{v} = -\frac{1}{\omega^2\rho'_{22}}\nabla\sigma - \frac{\rho'_{12}}{\rho'_{22}}\mathbf{u}. \quad (\text{C.24})$$

Substituting this into (C.20) and noting that

$$\epsilon = (\sigma - Qe)/R \quad (\text{C.25})$$

results in

$$\mu\Delta\mathbf{u} + \nabla\left(\left(A + \mu - \frac{Q^2}{R}\right)e\right) + \left(\frac{Q}{R} - \frac{\rho'_{12}}{\rho'_{22}}\right)\nabla\sigma + \frac{\rho'_{11}\rho'_{22} - (\rho'_{12})^2}{\rho'_{22}}\omega^2\mathbf{u} = 0. \quad (\text{C.26})$$

Now, taking the divergence of (C.21) gives

$$\Delta\sigma + \omega^2 (\rho'_{12}e + \rho'_{22}\epsilon) = 0, \quad (\text{C.27})$$

or, substituting in (C.25),

$$\Delta\sigma + \omega^2 \left(\rho'_{12}e + \rho'_{22}\frac{\sigma - Qe}{R}\right) = 0. \quad (\text{C.28})$$

Following [74], define the new variable

$$P = -\frac{\sigma}{\omega\sqrt{\rho'_{22}}}. \quad (\text{C.29})$$

Dividing (C.28) by $-\omega\sqrt{\rho'_{22}}$, in terms of the new variable P we obtain

$$\Delta P + \frac{\omega^2\rho_{22}}{R}P - \tilde{\alpha}e = 0, \quad (\text{C.30})$$

where the new parameter $\tilde{\alpha}$ is

$$\tilde{\alpha} = \sqrt{\rho'_{22}}\omega\left(\frac{\rho'_{12}}{\rho'_{22}} - \frac{Q}{R}\right). \quad (\text{C.31})$$

With this, the $\nabla\sigma$ term in (C.26) is merely $\tilde{\alpha}\nabla P$, so (C.26) can be written as

$$\mu\Delta\mathbf{u} + \nabla(\tilde{\lambda}e) + \tilde{\alpha}\nabla P + \tilde{\rho}\omega^2\mathbf{u} = 0, \quad (\text{C.32})$$

where

$$\begin{aligned} \tilde{\lambda} &= A + \mu - \frac{Q^2}{R} \\ \tilde{\rho} &= \frac{\rho'_{11}\rho'_{22} - (\rho'_{12})^2}{\rho'_{22}}. \end{aligned} \quad (\text{C.33})$$

Therefore, Biot's equations in terms of \mathbf{u} and P are

$$\mu \Delta \mathbf{u} + \nabla(\tilde{\lambda}e) + \tilde{\alpha}\nabla P + \tilde{\rho}\omega^2\mathbf{u} = 0 \quad (\text{C.34})$$

$$\Delta P + \omega^2\beta P - \tilde{\alpha}e = 0, \quad (\text{C.35})$$

where $\beta = \rho'_{22}/R$.

The fluid displacement \mathbf{v} can later be recovered from \mathbf{u} and P using the relation

$$\mathbf{v} = \frac{\nabla P}{\omega\sqrt{\rho'_{22}}} - \frac{\rho'_{12}}{\rho'_{22}}\mathbf{u}. \quad (\text{C.36})$$

C.4 Application of the Helmholtz decomposition to Biot's equations

Following [74], let

$$\mathbf{u} = \nabla\phi_f + \nabla\phi_s + \text{curl } \boldsymbol{\psi}, \quad \text{div } \boldsymbol{\psi} = 0 \quad (\text{C.37})$$

$$P = -\frac{\sigma}{\sqrt{\rho'_{22}\omega}} = A_f \Delta \phi_f + A_s \Delta \phi_s, \quad (\text{C.38})$$

where ϕ_f is the fast P wave potential, ϕ_s is the slow P wave potential, and $\boldsymbol{\psi}$ is the shear wave potential. The constants A_f and A_s will be determined later from a compatibility condition similar to that of [74].

The goal now is to obtain three Helmholtz equations for the three potentials ϕ_f , ϕ_s , and $\boldsymbol{\psi}$. Substituting (C.37) and (C.38) into (C.34) and (C.35), we get

$$\begin{aligned} \nabla \left\{ (\mu \Delta + \tilde{\lambda} \Delta + \tilde{\alpha} A_f \Delta + \tilde{\rho}\omega^2) \phi_f + (\mu \Delta + \tilde{\lambda} \Delta + \tilde{\alpha} A_s \Delta + \tilde{\rho}\omega^2) \phi_s \right\} \\ + \text{curl} \left\{ (\mu \Delta + \tilde{\rho}\omega^2) \boldsymbol{\psi} \right\} = 0 \end{aligned} \quad (\text{C.39})$$

$$\Delta \left\{ (A_f \Delta + \beta\omega^2 A_f - \tilde{\alpha}) \phi_f + (A_s \Delta + \beta\omega^2 A_s - \tilde{\alpha}) \phi_s \right\} = 0. \quad (\text{C.40})$$

If we can find solutions to the equations that come from setting everything in the curly braces to zero, then the problem is solved. The equations to be solved are now

$$(\mu + \tilde{\lambda} + \tilde{\alpha} A_f) \Delta \phi_f + \tilde{\rho}\omega^2 \phi_f + (\mu + \tilde{\lambda} + \tilde{\alpha} A_s) \Delta \phi_s + \tilde{\rho}\omega^2 \phi_s = 0 \quad (\text{C.41})$$

$$\mu \Delta \boldsymbol{\psi} + \tilde{\rho}\omega^2 \boldsymbol{\psi} = 0 \quad (\text{C.42})$$

$$A_f \Delta \phi_f + (\beta\omega^2 A_f - \tilde{\alpha}) \phi_f + A_s \Delta \phi_s + (\beta\omega^2 A_s - \tilde{\alpha}) \phi_s = 0. \quad (\text{C.43})$$

Equation (C.42) is already in the form of a Helmholtz equation.

Because this is a two-dimensional plane strain problem, $\boldsymbol{\psi}$ must point in the z direction (taking the problem to be stated in the x - y plane). Working in cylindrical coordinates, the r and θ components of $\boldsymbol{\psi}$ are both zero, so the z component will be labeled simply as $\psi(r, \theta)$; (C.42) then reduces to

$$\Delta \psi + k_t^2 \psi = 0, \quad (\text{C.44})$$

where

$$k_t^2 = \frac{\tilde{\rho}\omega^2}{\mu}, \quad \text{Im } k_t \geq 0. \quad (\text{C.45})$$

From (C.41) and (C.43), following [74], define

$$\begin{aligned}
a_f &= \mu + \tilde{\lambda} + \tilde{\alpha}A_f \\
a_s &= \mu + \tilde{\lambda} + \tilde{\alpha}A_s \\
d_f &= \beta\omega^2 A_f - \tilde{\alpha} \\
d_s &= \beta\omega^2 A_s - \tilde{\alpha}.
\end{aligned} \tag{C.46}$$

Then (C.41) and (C.43) can be written more compactly as

$$\begin{aligned}
a_f \Delta \phi_f + \tilde{\rho}\omega^2 \phi_f + a_s \Delta \phi_s + \tilde{\rho}\omega^2 \phi_s &= 0 \\
A_f \Delta \phi_f + d_f \phi_f + A_s \Delta \phi_s + d_s \phi_s &= 0.
\end{aligned} \tag{C.47}$$

For these equations to hold simultaneously, the compatibility condition in [74] implies that

$$\frac{d_f}{A_f} = \frac{\tilde{\rho}\omega^2}{a_f} \quad \text{and} \quad \frac{d_s}{A_s} = \frac{\tilde{\rho}\omega^2}{a_s}. \tag{C.48}$$

Choosing the expression involving A_f and expanding d_f and a_f back into their definitions from A_f results in

$$\frac{\beta\omega^2 A_f - \tilde{\alpha}}{A_f} = \frac{\tilde{\rho}\omega^2}{\mu + \tilde{\lambda} + \tilde{\alpha}A_f}, \tag{C.49}$$

or, multiplying out,

$$\tilde{\alpha}\beta\omega^2 A_f^2 + (\beta\omega^2(\mu + \tilde{\lambda}) - \tilde{\alpha}^2 - \tilde{\rho}\omega^2)A_f - \tilde{\alpha}(\mu + \tilde{\lambda}) = 0. \tag{C.50}$$

Letting $D = \beta\omega^2(\mu + \tilde{\lambda}) - \tilde{\alpha}^2 - \tilde{\rho}\omega^2$, the roots of this quadratic are

$$A_f = \frac{-D \pm \sqrt{D^2 + 4\tilde{\alpha}^2\beta\omega^2(\mu + \tilde{\lambda})}}{2\tilde{\alpha}\beta\omega^2}. \tag{C.51}$$

A_s satisfies the same quadratic equation, so it will be the opposite root from A_f . This results in the separate equations for ϕ_f and ϕ_s ,

$$\Delta \phi_f + k_f^2 \phi_f = 0, \quad k_f^2 = \frac{d_f}{A_f}, \quad \text{Im } k_f \geq 0 \tag{C.52}$$

$$\Delta \phi_s + k_s^2 \phi_s = 0, \quad k_s^2 = \frac{d_s}{A_s}, \quad \text{Im } k_s \geq 0. \tag{C.53}$$

A few notes on these Helmholtz equations are in order. First, the condition $\text{Im } k \geq 0$ ensures that the wave decays as it propagates outward. Second, A_f is chosen so that $\text{Re } k_f \leq \text{Re } k_s$, ensuring that ϕ_f really corresponds to the fast P wave

One last PDE must be stated before solution can begin: the Helmholtz equation describing the acoustic pressure p_w in the surrounding fluid,

$$\Delta p_w + \frac{\omega^2 \rho_0}{K_0} p_w = 0, \tag{C.54}$$

where the pressure field in the fluid is $p_w(\mathbf{x})e^{-i\omega t}$, ρ_0 is the density of the fluid, and K_0 is its bulk modulus.

C.5 Solution of the Helmholtz equations for the wave potentials

C.5.1 Problem setting

To summarize, the problem of finding the state vector $\mathbf{Q}(x, z, t)$ for a poroelastic cylinder scattering a time-harmonic incoming acoustic wave has been transformed to one of solving a system of three decoupled scalar Helmholtz equations describing the poroelastic waves within the disc, and one scalar Helmholtz equation describing the waves in the fluid. Let the origin of coordinates be at the disc center, and let the radius of the disc be a . Then in cylindrical coordinates the PDEs describing the system are

$$\left. \begin{aligned} \Delta \phi_f + k_f^2 \phi_f &= 0 \\ \Delta \phi_s + k_s^2 \phi_s &= 0 \\ \Delta \psi + k_t^2 \psi &= 0 \end{aligned} \right\}, \quad r < a \quad (\text{C.55})$$

$$\Delta p_w + \frac{\omega^2 \rho_0}{K_0} p_w = 0, \quad r > a \quad (\text{C.56})$$

together with a set of transmission conditions at $r = a$:

$$\left(\mathbf{U}^{(i)} + \mathbf{U}^{(r)} \right) \cdot \mathbf{n} = \left(\mathbf{u}^{(f)} + \phi(\mathbf{U}^{(f)} - \mathbf{u}^{(f)}) \right) \cdot \mathbf{n} \quad (\text{conservation of fluid mass}) \quad (\text{C.57})$$

$$\eta_d \left(p_f^{(f)} - (p_w^{(i)} + p_w^{(r)}) \right) = Z(1 - \eta_d) \phi \left(\dot{\mathbf{U}}^{(f)} - \dot{\mathbf{u}}^{(f)} \right) \cdot \mathbf{n} \quad (\text{hydraulic contact condition}) \quad (\text{C.58})$$

$$- \left(p_w^{(i)} + p_w^{(r)} \right) = \tau_{rr}^{(f)} \quad (\text{normal traction continuity}) \quad (\text{C.59})$$

$$0 = \tau_{r\theta}^{(f)} \quad (\text{tangential traction continuity}). \quad (\text{C.60})$$

Here the superscript (i) indicates the incident wave, (r) indicates the reflected wave, and (f) indicates the refracted waves; Z is the fluid acoustic impedance, $Z = \sqrt{K_0 \rho_0}$, and \mathbf{n} is the unit normal vector pointing out of the poroelastic medium. The parameter η_d is the interface discharge efficiency between the fluid and the poroelastic medium; $\eta_d = 1$ models open pores, $\eta_d = 0$ models sealed pores, and values between 0 and 1 indicate imperfect hydraulic contact. Transforming into the frequency domain and simplifying a bit,

the transmission conditions become

$$\left(\mathbf{v}_w^{(i)} + \mathbf{v}_w^{(r)}\right) \cdot \mathbf{n} = \left((1 - \phi)\mathbf{u}^{(f)} + \phi\mathbf{v}^{(f)}\right) \cdot \mathbf{n} \quad (\text{conservation of fluid mass}) \quad (\text{C.61})$$

$$\eta_d \left(p_f^{(f)} - (p_w^{(i)} + p_w^{(r)})\right) = -i\omega\phi Z(1 - \eta_d) \left(\mathbf{v}^{(f)} - \mathbf{u}^{(f)}\right) \cdot \mathbf{n} \quad (\text{hydraulic contact condition}) \quad (\text{C.62})$$

$$- \left(p_w^{(i)} + p_w^{(r)}\right) = \tau_{rr}^{(f)} \quad (\text{normal traction continuity}) \quad (\text{C.63})$$

$$0 = \tau_{r\theta}^{(f)} \quad (\text{tangential traction continuity}). \quad (\text{C.64})$$

We specify the spatial dependence of the incident wave as

$$p_w^{(i)} = p_0 e^{ikx} = p_0 e^{ikr \cos \theta}, \quad (\text{C.65})$$

where k is

$$k = \omega \sqrt{\frac{\rho_0}{K_0}}. \quad (\text{C.66})$$

C.5.2 Transmission conditions in terms of potentials

Before moving on to the functional form of the solution, it will be useful to first express the transmission conditions (C.61) through (C.64) in terms of the potentials ϕ_f , ϕ_s , ψ , and p_w .

First, from (C.37), the solid displacement $\mathbf{u}^{(f)}$ can be expressed in terms of ϕ_f , ϕ_s , and ψ as

$$\mathbf{u}^{(f)} = \nabla\phi_f + \nabla\phi_s + \text{curl } \boldsymbol{\psi}. \quad (\text{C.67})$$

In polar coordinates, the gradient operator is

$$\nabla = \mathbf{e}_r \frac{\partial}{\partial r} + \mathbf{e}_\theta \frac{1}{r} \frac{\partial}{\partial \theta}, \quad (\text{C.68})$$

and the curl of $\boldsymbol{\psi} = \psi \mathbf{e}_z$ is

$$\text{curl } \boldsymbol{\psi} = \frac{1}{r} \frac{\partial \psi}{\partial \theta} \mathbf{e}_r - \frac{\partial \psi}{\partial r} \mathbf{e}_\theta, \quad (\text{C.69})$$

so $\mathbf{u}^{(f)}$ can be expressed explicitly in terms of the potentials as

$$\mathbf{u}^{(f)} = \left(\frac{\partial \phi_f}{\partial r} + \frac{\partial \phi_s}{\partial r} + \frac{1}{r} \frac{\partial \psi}{\partial \theta}\right) \mathbf{e}_r + \left(\frac{1}{r} \frac{\partial \phi_f}{\partial \theta} + \frac{1}{r} \frac{\partial \phi_s}{\partial \theta} - \frac{\partial \psi}{\partial r}\right) \mathbf{e}_\theta. \quad (\text{C.70})$$

Now, combining (C.36) and (C.38), the fluid displacement $\mathbf{v}^{(f)}$ can be expressed in terms of the potentials ϕ_f and ϕ_s , and the solid displacement $\mathbf{u}^{(f)}$, as

$$\mathbf{v}^{(f)} = \frac{1}{\omega \sqrt{\rho'_{22}}} (A_f \nabla \Delta \phi_f + A_s \nabla \Delta \phi_s) - \frac{\rho'_{12}}{\rho'_{22}} \mathbf{u}^{(f)}. \quad (\text{C.71})$$

Remembering that $\Delta \phi_{f,s} + k_{f,s}^2 \phi_{f,s} = 0$, as well as the relation (C.37), this reduces to

$$\mathbf{v}^{(f)} = - \left(\frac{k_f^2 A_f}{\omega \sqrt{\rho'_{22}}} + \frac{\rho'_{12}}{\rho'_{22}} \right) \nabla \phi_f - \left(\frac{k_s^2 A_s}{\omega \sqrt{\rho'_{22}}} + \frac{\rho'_{12}}{\rho'_{22}} \right) \nabla \phi_s - \frac{\rho'_{12}}{\rho'_{22}} \text{curl } \boldsymbol{\psi}. \quad (\text{C.72})$$

Expanding this in terms of the gradient and curl in cylindrical coordinates gives

$$\begin{aligned} \mathbf{v}^{(f)} = & - \left(\left(\frac{k_f^2 A_f}{\omega \sqrt{\rho'_{22}}} + \frac{\rho'_{12}}{\rho'_{22}} \right) \frac{\partial \phi_f}{\partial r} + \left(\frac{k_s^2 A_s}{\omega \sqrt{\rho'_{22}}} + \frac{\rho'_{12}}{\rho'_{22}} \right) \frac{\partial \phi_s}{\partial r} + \frac{\rho'_{12}}{\rho'_{22} r} \frac{\partial \psi}{\partial \theta} \right) \mathbf{e}_r \\ & - \left(\left(\frac{k_f^2 A_f}{\omega \sqrt{\rho'_{22}}} + \frac{\rho'_{12}}{\rho'_{22}} \right) \frac{1}{r} \frac{\partial \phi_f}{\partial \theta} + \left(\frac{k_s^2 A_s}{\omega \sqrt{\rho'_{22}}} + \frac{\rho'_{12}}{\rho'_{22}} \right) \frac{1}{r} \frac{\partial \phi_s}{\partial \theta} - \frac{\rho'_{12}}{\rho'_{22}} \frac{\partial \psi}{\partial r} \right) \mathbf{e}_\theta. \end{aligned} \quad (\text{C.73})$$

For the fluid displacement outside the poroelastic disc, first note that the fluid acceleration $\ddot{\mathbf{U}}_w$ is related to the pressure gradient by

$$\rho_0 \ddot{\mathbf{U}}_w = -\nabla p_w. \quad (\text{C.74})$$

For the time-harmonic case, this becomes

$$-\omega^2 \rho_0 \mathbf{v}_w = -\nabla p_w, \quad (\text{C.75})$$

so the displacement of the fluid outside the disc is

$$\mathbf{v}_w = \frac{1}{\omega^2 \rho_0} \nabla p_w = \frac{1}{\omega^2 \rho_0} \frac{\partial p_w}{\partial r} \mathbf{e}_r + \frac{1}{\omega^2 \rho_0 r} \frac{\partial p_w}{\partial \theta} \mathbf{e}_\theta. \quad (\text{C.76})$$

Since the unit outward normal \mathbf{n} of the disc is just $\mathbf{n} = \mathbf{e}_r$, putting $r = a$, the first transmission condition (C.61) becomes

$$\begin{aligned} \frac{1}{\omega^2 \rho_0} \left(\frac{\partial p_w^{(i)}}{\partial r} + \frac{\partial p_w^{(r)}}{\partial r} \right) &= (1 - \phi) \left(\frac{\partial \phi_f}{\partial r} + \frac{\partial \phi_s}{\partial r} + \frac{1}{a} \frac{\partial \psi}{\partial \theta} \right) \\ &\quad - \phi \left(\left(\frac{k_f^2 A_f}{\omega \sqrt{\rho'_{22}}} + \frac{\rho'_{12}}{\rho'_{22}} \right) \frac{\partial \phi_f}{\partial r} + \left(\frac{k_s^2 A_s}{\omega \sqrt{\rho'_{22}}} + \frac{\rho'_{12}}{\rho'_{22}} \right) \frac{\partial \phi_s}{\partial r} + \frac{\rho'_{12}}{\rho'_{22} a} \frac{\partial \psi}{\partial \theta} \right) \\ &= \delta_{1f} \frac{\partial \phi_f}{\partial r} + \delta_{1s} \frac{\partial \phi_s}{\partial r} + \delta_{1t} \frac{\partial \psi}{\partial \theta}, \end{aligned} \quad (\text{C.77})$$

where the new δ coefficients are

$$\begin{aligned} \delta_{1f} &= 1 - \phi \left(1 + \frac{k_f^2 A_f}{\omega \sqrt{\rho'_{22}}} + \frac{\rho'_{12}}{\rho'_{22}} \right) \\ \delta_{1s} &= 1 - \phi \left(1 + \frac{k_s^2 A_s}{\omega \sqrt{\rho'_{22}}} + \frac{\rho'_{12}}{\rho'_{22}} \right) \\ \delta_{1t} &= \frac{1}{a} \left(1 - \phi \left(1 + \frac{\rho'_{12}}{\rho'_{22}} \right) \right). \end{aligned} \quad (\text{C.78})$$

The notation for these coefficients will be as follows: the Greek letter will denote which order of derivative they multiply (γ for the potentials themselves, δ for first derivatives, etc.), the first subscript will denote which transmission condition they correspond to, and the second subscript will denote which potential they multiply.

Writing the next transmission condition in terms of the potentials will require the fluid pressure in the poroelastic medium, p_f . From (C.38), (C.29), (C.18), and the fact that $\Delta \phi_{f,s} + k_{f,s}^2 \phi_{f,s} = 0$, we get

$$p_f = \frac{\omega \sqrt{\rho'_{22}}}{\phi} P = -\frac{\omega \sqrt{\rho'_{22}}}{\phi} (k_f^2 A_f \phi_f + k_s^2 A_s \phi_s) \quad (\text{C.79})$$

Thus the hydraulic contact condition (C.62) becomes

$$\begin{aligned} \eta_d \left(-\frac{\omega \sqrt{\rho'_{22}}}{\phi} (k_f^2 A_f \phi_f + k_s^2 A_s \phi_s) - (p_w^{(i)} + p_w^{(r)}) \right) &= i\omega \phi Z (1 - \eta_d) \\ \cdot \left(\left(\frac{k_f^2 A_f}{\omega \sqrt{\rho'_{22}}} + \frac{\rho'_{12}}{\rho'_{22}} \right) \frac{\partial \phi_f}{\partial r} + \left(\frac{k_s^2 A_s}{\omega \sqrt{\rho'_{22}}} + \frac{\rho'_{12}}{\rho'_{22}} \right) \frac{\partial \phi_s}{\partial r} + \frac{\rho'_{12}}{\rho'_{22} r} \frac{\partial \psi}{\partial \theta} + \left(\frac{\partial \phi_f}{\partial r} + \frac{\partial \phi_s}{\partial r} + \frac{1}{r} \frac{\partial \psi}{\partial \theta} \right) \right). \end{aligned} \quad (\text{C.80})$$

Simplifying yields

$$-\eta_d (p_w^{(i)} + p_w^{(r)}) = \delta_{2f} \frac{\partial \phi_f}{\partial r} + \gamma_{2f} \phi_f + \delta_{2s} \frac{\partial \phi_s}{\partial r} + \gamma_{2s} \phi_s + \delta_{2t} \frac{\partial \psi}{\partial \theta}, \quad (\text{C.81})$$

where

$$\begin{aligned} \gamma_{2f} &= \frac{\eta_d \omega \sqrt{\rho'_{22}} k_f^2 A_f}{\phi} \\ \gamma_{2s} &= \frac{\eta_d \omega \sqrt{\rho'_{22}} k_s^2 A_s}{\phi} \\ \delta_{2f} &= i\omega \phi Z (1 - \eta_d) \left(1 + \frac{k_f^2 A_f}{\omega \sqrt{\rho'_{22}}} + \frac{\rho'_{12}}{\rho'_{22}} \right) \\ \delta_{2s} &= i\omega \phi Z (1 - \eta_d) \left(1 + \frac{k_s^2 A_s}{\omega \sqrt{\rho'_{22}}} + \frac{\rho'_{12}}{\rho'_{22}} \right) \\ \delta_{2t} &= \frac{i\omega \phi Z (1 - \eta_d)}{a} \left(1 + \frac{\rho'_{12}}{\rho'_{22}} \right). \end{aligned} \quad (\text{C.82})$$

Formulating the remaining transmission conditions will require the total stress tensor $\boldsymbol{\tau}$ in the poroelastic medium. From (C.7), (C.13) and the definition $\sigma = -\phi p_f$, we get

$$\tau_{ij} = \sigma_{ij} + \delta_{ij} \sigma = 2N e_{ij} + \delta_{ij} (Ae + Q\epsilon - \phi p_f). \quad (\text{C.83})$$

The fluid pressure in the poroelastic medium p_f is already known in terms of the potentials from (C.79), so only the strains e_{ij} and the dilatations e and ϵ are still needed. First, the dilatations are

$$e = \text{div } \mathbf{u}^{(f)} = \text{div}(\nabla \phi_f + \nabla \phi_s + \text{curl } \boldsymbol{\psi}) = \Delta \phi_f + \Delta \phi_s = -(k_f^2 \phi_f + k_s^2 \phi_s) \quad (\text{C.84})$$

and, via (C.72)

$$\begin{aligned}\epsilon = \operatorname{div} \mathbf{v}^{(f)} &= \operatorname{div} \left(- \left(\frac{k_f^2 A_f}{\omega \sqrt{\rho'_{22}}} + \frac{\rho'_{12}}{\rho'_{22}} \right) \nabla \phi_f - \left(\frac{k_s^2 A_s}{\omega \sqrt{\rho'_{22}}} + \frac{\rho'_{12}}{\rho'_{22}} \right) \nabla \phi_s - \frac{\rho'_{12}}{\rho'_{22}} \operatorname{curl} \psi \right) \\ &= \left(\frac{k_f^2 A_f}{\omega \sqrt{\rho'_{22}}} + \frac{\rho'_{12}}{\rho'_{22}} \right) k_f^2 \phi_f + \left(\frac{k_s^2 A_s}{\omega \sqrt{\rho'_{22}}} + \frac{\rho'_{12}}{\rho'_{22}} \right) k_s^2 \phi_s. \quad (\text{C.85})\end{aligned}$$

Now, from the relations between strain and displacement in polar coordinates, the tensor strains e_{ij} are

$$e_{rr} = \frac{\partial u_r^{(f)}}{\partial r} = \frac{\partial^2 \phi_f}{\partial r^2} + \frac{\partial^2 \phi_s}{\partial r^2} - \frac{1}{r^2} \frac{\partial \psi}{\partial \theta} + \frac{1}{r} \frac{\partial^2 \psi}{\partial \theta \partial r} \quad (\text{C.86})$$

$$e_{\theta\theta} = \frac{1}{r} \frac{\partial u_\theta^{(f)}}{\partial \theta} + \frac{u_r^{(f)}}{r} = \frac{1}{r^2} \left(\frac{\partial^2 \phi_f}{\partial \theta^2} + \frac{\partial^2 \phi_s}{\partial \theta^2} + \frac{\partial \psi}{\partial \theta} \right) + \frac{1}{r} \left(\frac{\partial \phi_f}{\partial r} + \frac{\partial \phi_s}{\partial r} - \frac{\partial^2 \psi}{\partial \theta \partial r} \right) \quad (\text{C.87})$$

$$\begin{aligned}e_{r\theta} &= \frac{1}{2} \left(\frac{1}{r} \frac{\partial u_r^{(f)}}{\partial \theta} + \frac{\partial u_\theta^{(f)}}{\partial r} - \frac{u_\theta^{(f)}}{r} \right) \\ &= \frac{1}{2} \left(\frac{1}{r^2} \left(\frac{\partial^2 \psi}{\partial \theta^2} - 2 \frac{\partial \phi_f}{\partial \theta} - 2 \frac{\partial \phi_s}{\partial \theta} \right) + \frac{1}{r} \left(2 \frac{\partial^2 \phi_f}{\partial \theta \partial r} + 2 \frac{\partial^2 \phi_s}{\partial \theta \partial r} + \frac{\partial \psi}{\partial r} \right) - \frac{\partial^2 \psi}{\partial r^2} \right). \quad (\text{C.88})\end{aligned}$$

Substituting these expressions, along with (C.79), into (C.83) gives the total stress components directly in terms of the potentials:

$$\begin{aligned}\tau_{rr} &= 2N e_{rr} + A e + Q \epsilon - \phi p_f \\ &= 2N \left(\frac{\partial^2 \phi_f}{\partial r^2} + \frac{\partial^2 \phi_s}{\partial r^2} - \frac{1}{r^2} \frac{\partial \psi}{\partial \theta} + \frac{1}{r} \frac{\partial^2 \psi}{\partial \theta \partial r} \right) - A (k_f^2 \phi_f + k_s^2 \phi_s) \\ &\quad + Q \left(\left(\frac{k_f^2 A_f}{\omega \sqrt{\rho'_{22}}} + \frac{\rho'_{12}}{\rho'_{22}} \right) k_f^2 \phi_f + \left(\frac{k_s^2 A_s}{\omega \sqrt{\rho'_{22}}} + \frac{\rho'_{12}}{\rho'_{22}} \right) k_s^2 \phi_s \right) \\ &\quad + \omega \sqrt{\rho'_{22}} (k_f^2 A_f \phi_f + k_s^2 A_s \phi_s) \\ &= 2N \left(\frac{\partial^2 \phi_f}{\partial r^2} + \frac{\partial^2 \phi_s}{\partial r^2} - \frac{1}{r^2} \frac{\partial \psi}{\partial \theta} + \frac{1}{r} \frac{\partial^2 \psi}{\partial \theta \partial r} \right) \\ &\quad + \left(-A + Q \left(\frac{k_f^2 A_f}{\omega \sqrt{\rho'_{22}}} + \frac{\rho'_{12}}{\rho'_{22}} \right) + \omega \sqrt{\rho'_{22}} A_f \right) k_f^2 \phi_f \\ &\quad + \left(-A + Q \left(\frac{k_s^2 A_s}{\omega \sqrt{\rho'_{22}}} + \frac{\rho'_{12}}{\rho'_{22}} \right) + \omega \sqrt{\rho'_{22}} A_s \right) k_s^2 \phi_s\end{aligned} \quad (\text{C.89})$$

$$\begin{aligned}
\tau_{\theta\theta} &= 2Ne_{\theta\theta} + Ae + Q\epsilon - \phi p_f \\
&= 2N \left(\frac{1}{r^2} \left(\frac{\partial^2 \phi_f}{\partial \theta^2} + \frac{\partial^2 \phi_s}{\partial \theta^2} + \frac{\partial \psi}{\partial \theta} \right) + \frac{1}{r} \left(\frac{\partial \phi_f}{\partial r} + \frac{\partial \phi_s}{\partial r} - \frac{\partial^2 \psi}{\partial \theta \partial r} \right) \right) \\
&\quad + \left(-A + Q \left(\frac{k_f^2 A_f}{\omega \sqrt{\rho'_{22}}} + \frac{\rho'_{12}}{\rho'_{22}} \right) + \omega \sqrt{\rho'_{22}} A_f \right) k_f^2 \phi_f \\
&\quad + \left(-A + Q \left(\frac{k_s^2 A_s}{\omega \sqrt{\rho'_{22}}} + \frac{\rho'_{12}}{\rho'_{22}} \right) + \omega \sqrt{\rho'_{22}} A_s \right) k_s^2 \phi_s
\end{aligned} \tag{C.90}$$

$$\tau_{r\theta} = 2Ne_{r\theta} = N \left(\frac{1}{r^2} \left(\frac{\partial^2 \psi}{\partial \theta^2} - 2 \frac{\partial \phi_f}{\partial \theta} - 2 \frac{\partial \phi_s}{\partial \theta} \right) + \frac{1}{r} \left(2 \frac{\partial^2 \phi_f}{\partial \theta \partial r} + 2 \frac{\partial^2 \phi_s}{\partial \theta \partial r} + \frac{\partial \psi}{\partial r} \right) - \frac{\partial^2 \psi}{\partial r^2} \right) \tag{C.91}$$

Transmission condition (C.63) then becomes

$$- (p_w^{(i)} + p_w^{(r)}) = \epsilon_{3f} \frac{\partial^2 \phi_f}{\partial r^2} + \epsilon_{3s} \frac{\partial^2 \phi_s}{\partial r^2} + \delta_{3t} \frac{\partial \psi}{\partial \theta} + \epsilon_{3t} \frac{\partial^2 \psi}{\partial \theta \partial r} + \gamma_{3f} \phi_f + \gamma_{3s} \phi_s, \tag{C.92}$$

where the coefficients are

$$\begin{aligned}
\gamma_{3f} &= \left(-A + Q \left(\frac{k_f^2 A_f}{\omega \sqrt{\rho'_{22}}} + \frac{\rho'_{12}}{\rho'_{22}} \right) + \omega \sqrt{\rho'_{22}} A_f \right) k_f^2 \\
\gamma_{3s} &= \left(-A + Q \left(\frac{k_s^2 A_s}{\omega \sqrt{\rho'_{22}}} + \frac{\rho'_{12}}{\rho'_{22}} \right) + \omega \sqrt{\rho'_{22}} A_s \right) k_s^2 \\
\delta_{3t} &= -\frac{2N}{a^2} \\
\epsilon_{3f} &= 2N \\
\epsilon_{3s} &= 2N \\
\epsilon_{3t} &= \frac{2N}{a}.
\end{aligned} \tag{C.93}$$

Similarly, transmission condition (C.64) becomes

$$0 = \delta_{4f} \frac{\partial \phi_f}{\partial \theta} + \delta_{4s} \frac{\partial \phi_s}{\partial \theta} + \delta_{4t} \frac{\partial \psi}{\partial r} + \epsilon_{4f} \frac{\partial^2 \phi_f}{\partial \theta \partial r} + \epsilon_{4s} \frac{\partial^2 \phi_s}{\partial \theta \partial r} + \epsilon_{4tr} \frac{\partial^2 \psi}{\partial r^2} + \epsilon_{4t\theta} \frac{\partial^2 \psi}{\partial \theta^2} \tag{C.94}$$

with

$$\begin{aligned}
\delta_{4f} &= -\frac{2}{a^2} \\
\delta_{4s} &= -\frac{2}{a^2} \\
\delta_{4t} &= \frac{1}{a} \\
\epsilon_{4f} &= \frac{2}{a} \\
\epsilon_{4s} &= \frac{2}{a} \\
\epsilon_{4tr} &= -1 \\
\epsilon_{4t\theta} &= \frac{1}{a^2}
\end{aligned} \tag{C.95}$$

Note that a common factor of N has been divided out in forming this equation.

In summary, in terms of the potentials, the transmission conditions are

$$\frac{1}{\omega^2 \rho_0} \left(\frac{\partial p_w^{(i)}}{\partial r} + \frac{\partial p_w^{(r)}}{\partial r} \right) = \delta_{1f} \frac{\partial \phi_f}{\partial r} + \delta_{1s} \frac{\partial \phi_s}{\partial r} + \delta_{1t} \frac{\partial \psi}{\partial \theta} \tag{C.96}$$

$$-\eta_d(p_w^{(i)} + p_w^{(r)}) = \delta_{2f} \frac{\partial \phi_f}{\partial r} + \gamma_{2f} \phi_f + \delta_{2s} \frac{\partial \phi_s}{\partial r} + \gamma_{2s} \phi_s + \delta_{2t} \frac{\partial \psi}{\partial \theta} \tag{C.97}$$

$$-(p_w^{(i)} + p_w^{(r)}) = \epsilon_{3f} \frac{\partial^2 \phi_f}{\partial r^2} + \epsilon_{3s} \frac{\partial^2 \phi_s}{\partial r^2} + \delta_{3t} \frac{\partial \psi}{\partial \theta} + \epsilon_{3t} \frac{\partial^2 \psi}{\partial \theta \partial r} + \gamma_{3f} \phi_f + \gamma_{3s} \phi_s \tag{C.98}$$

$$\begin{aligned}
0 &= \delta_{4f} \frac{\partial \phi_f}{\partial \theta} + \delta_{4s} \frac{\partial \phi_s}{\partial \theta} + \delta_{4t} \frac{\partial \psi}{\partial r} \\
&\quad + \epsilon_{4f} \frac{\partial^2 \phi_f}{\partial \theta \partial r} + \epsilon_{4s} \frac{\partial^2 \phi_s}{\partial \theta \partial r} + \epsilon_{4tr} \frac{\partial^2 \psi}{\partial r^2} + \epsilon_{4t\theta} \frac{\partial^2 \psi}{\partial \theta^2}
\end{aligned} \tag{C.99}$$

where the coefficients are given in (C.78), (C.82), (C.93), and (C.95).

C.5.3 Bessel function expansions

The reflected wave $p_w^{(r)}$ solves the Helmholtz equation outside the disc, so by the results of [64], it can be expressed as a series,

$$p_w^{(r)} = \sum_{m=-\infty}^{\infty} P_m H_m(kr) e^{im\theta}, \tag{C.100}$$

where H_m is the Hankel function of the first kind of order m . Similarly for the refracted wave potentials, since they satisfy Helmholtz equations inside the disc, from [64] they may

also be expressed as series,

$$\phi_f = \sum_{m=-\infty}^{\infty} F_m J_m(k_f r) e^{im\theta} \quad (\text{C.101})$$

$$\phi_s = \sum_{m=-\infty}^{\infty} S_m J_m(k_s r) e^{im\theta} \quad (\text{C.102})$$

$$\psi = \sum_{m=-\infty}^{\infty} T_m J_m(k_t r) e^{im\theta}, \quad (\text{C.103})$$

where J_m is the Bessel function of the first kind of order m .

In addition, by the identity on page 1371 of [64], the incident wave can also be expanded as a Bessel function series,

$$p_w^{(i)} = p_0 e^{ikr \cos \theta} = p_0 \sum_{m=-\infty}^{\infty} i^m J_m(kr) e^{im\theta}. \quad (\text{C.104})$$

C.5.4 Transmission conditions in terms of series coefficients

Substituting the Bessel function series expansions of the previous section into the transmission conditions (C.96) through (C.99), and noting the linear independence of the functions $e^{im\theta}$, $m \in \mathbb{Z}$, we can derive relations between the series coefficients P_m , F_m , S_m , and T_m for any m :

$$\frac{1}{\omega^2 \rho_0} (p_0 i^m k J'_m(ka) + P_m k H'_m(ka)) = \delta_{1f} F_m k_f J'_m(k_f a) + \delta_{1s} S_m k_s J'_m(k_s a) + \delta_{1t} i m T_m J_m(k_t a) \quad (\text{C.105})$$

$$\begin{aligned} -\eta_d (p_0 i^m J_m(ka) + P_m H_m(ka)) &= \delta_{2f} F_m k_f J'_m(k_f a) + \gamma_{2f} F_m J_m(k_f a) \\ &+ \delta_{2s} S_m k_s J'_m(k_s a) + \gamma_{2s} S_m J_m(k_s a) + \delta_{2t} i m T_m J_m(k_t a) \end{aligned} \quad (\text{C.106})$$

$$\begin{aligned} -(p_0 i^m J_m(ka) + P_m H_m(ka)) &= \epsilon_{3f} F_m k_f^2 J''_m(k_f a) + \epsilon_{3s} S_m k_s^2 J''_m(k_s a) + \delta_{3t} i m T_m J_m(k_t a) \\ &+ \epsilon_{3t} i m T_m k_t J'_m(k_t a) + \gamma_{3f} F_m J_m(k_f a) + \gamma_{3s} S_m J_m(k_s a) \end{aligned} \quad (\text{C.107})$$

$$\begin{aligned} 0 &= \delta_{4f} i m F_m J_m(k_f a) + \delta_{4s} i m S_m J_m(k_s a) + \delta_{4t} k_t T_m J'_m(k_t a) \\ &+ \epsilon_{4f} i m k_f F_m J'_m(k_f a) + \epsilon_{4s} i m k_s S_m J'_m(k_s a) + \epsilon_{4t} k_t^2 T_m J''_m(k_t a) - \epsilon_{4t\theta} m^2 T_m J_m(k_t a) \end{aligned} \quad (\text{C.108})$$

For practical computation, the derivatives of the Bessel and Hankel functions are computed using recurrence relations such as $2J'_m = J_{m-1} - J_{m+1}$, $2H'_m = H_{m-1} - H_{m+1}$, etc.

Conditions (C.105) through (C.108) can be cast in matrix form as

$$\mathbf{M}_m \mathbf{c}_m = \mathbf{r}_m \quad (\text{C.109})$$

where

$$\mathbf{c}_m = \begin{pmatrix} P_m \\ F_m \\ S_m \\ T_m \end{pmatrix}, \quad \mathbf{r}_m = p_0 i^m \begin{pmatrix} kJ'_m(ka)/(\omega^2 \rho_0) \\ -\eta_d J_m(ka) \\ -J_m(ka) \\ 0 \end{pmatrix} \quad (\text{C.110})$$

and the elements of the coefficient matrix are

$$\begin{aligned} (\mathbf{M}_m)_{11} &= -\frac{kH'_m(ka)}{\omega^2 \rho_0} \\ (\mathbf{M}_m)_{12} &= \delta_{1f} k_f J'_m(k_f a) \\ (\mathbf{M}_m)_{13} &= \delta_{1s} k_s J'_m(k_s a) \\ (\mathbf{M}_m)_{14} &= im \delta_{1t} J_m(k_t a) \\ (\mathbf{M}_m)_{21} &= \eta_d H_m(ka) \\ (\mathbf{M}_m)_{22} &= \delta_{2f} k_f J'_m(k_f a) + \gamma_{2f} J_m(k_f a) \\ (\mathbf{M}_m)_{23} &= \delta_{2s} k_s J'_m(k_s a) + \gamma_{2s} J_m(k_s a) \\ (\mathbf{M}_m)_{24} &= im \delta_{2t} J_m(k_t a) \\ (\mathbf{M}_m)_{31} &= H_m(ka) \\ (\mathbf{M}_m)_{32} &= \epsilon_{3f} k_f^2 J''_m(k_f a) + \gamma_{3f} J_m(k_f a) \\ (\mathbf{M}_m)_{33} &= \epsilon_{3s} k_s^2 J''_m(k_s a) + \gamma_{3s} J_m(k_s a) \\ (\mathbf{M}_m)_{34} &= im(\delta_{3t} J_m(k_t a) + \epsilon_{3t} k_t J'_m(k_t a)) \\ (\mathbf{M}_m)_{41} &= 0 \\ (\mathbf{M}_m)_{42} &= im(\delta_{4f} J_m(k_f a) + \epsilon_{4f} k_f J'_m(k_f a)) \\ (\mathbf{M}_m)_{43} &= im(\delta_{4s} J_m(k_s a) + \epsilon_{4s} k_s J'_m(k_s a)) \\ (\mathbf{M}_m)_{44} &= \delta_{4t} k_t J'_m(k_t a) + \epsilon_{4tr} k_t^2 J''_m(k_t a) - \epsilon_{4t\theta} m^2 J_m(k_t a). \end{aligned} \quad (\text{C.111})$$

As written, the matrix \mathbf{M}_m is often extremely badly conditioned, with condition numbers beyond 10^{16} . However, this may be regarded as a sort of “fake” ill-conditioning, since rescaling the system by dividing each column of \mathbf{M}_m by the maximum magnitude of the elements of that column, then dividing each row by the maximum magnitude of each row, results in much more modest condition numbers, typically no worse than 10^4 or 10^5 for the cylindrical scatterer cases of Section 7.1. In practice, system (C.109) is solved by first rescaling it in this fashion, solving the rescaled linear system, then transforming back into the original variables.

C.6 Implementation of the solution

C.6.1 Variables of the first-order system in terms of the series representation

The Bessel function series expansion can also be used to write series representations of the variables of the first-order system used in the CLAWPACK code — the fluid pressure, total stress, solid velocity, and fluid flow rate.

First, starting from (C.79) and inserting a factor of $e^{-i\omega t}$ to account for time dependence, $p_f(r, \theta, t)$ can be obtained as

$$\begin{aligned} p_f(r, \theta, t) &= -e^{-i\omega t} \frac{\omega \sqrt{\rho'_{22}}}{\phi} (k_f^2 A_f \phi_f + k_s^2 A_s \phi_s) \\ &= \sum_{m=-\infty}^{\infty} \left((p_f)_{00}^{(f)} F_m J_m(k_f r) + (p_f)_{00}^{(s)} S_m J_m(k_s r) \right) e^{i(m\theta - \omega t)}, \end{aligned} \quad (\text{C.112})$$

where the coefficients $(p_f)_{00}^{(f)}$ and $(p_f)_{00}^{(s)}$ are

$$\begin{aligned} (p_f)_{00}^{(f)} &= -\frac{\omega \sqrt{\rho'_{22}}}{\phi} k_f^2 A_f \\ (p_f)_{00}^{(s)} &= -\frac{\omega \sqrt{\rho'_{22}}}{\phi} k_s^2 A_s. \end{aligned} \quad (\text{C.113})$$

Here the superscript in the coefficients denotes which wave they are associated with, the first character of the subscript is the order of the derivative of the associated Bessel function, and the second is the power of m present in the term. This nomenclature anticipates the writing of computer code for evaluating the solution, where it will be convenient to compute the series coefficients F_m , etc., and a small set of coefficients such as $(p_f)_{00}^{(f)}$ that are the same for all terms.

Next, substituting the Bessel function series expansions into equations (C.89), (C.90), and (C.91), and multiplying by $e^{-i\omega t}$, we arrive at series expansions for the total stress components in polar coordinates:

$$\begin{aligned} \tau_{rr} &= \sum_{m=-\infty}^{\infty} \left((\tau_{rr})_{00}^{(f)} F_m J_m(k_f r) + (\tau_{rr})_{20}^{(f)} F_m J_m''(k_f r) + (\tau_{rr})_{00}^{(s)} S_m J_m(k_s r) \right. \\ &\quad \left. + (\tau_{rr})_{20}^{(s)} S_m J_m''(k_s r) + (\tau_{rr})_{01}^{(t)} m T_m \frac{J_m(k_t r)}{r^2} + (\tau_{rr})_{11}^{(t)} m T_m \frac{J_m'(k_t r)}{r} \right) e^{i(m\theta - \omega t)} \end{aligned} \quad (\text{C.114})$$

$$\begin{aligned} \tau_{\theta\theta} &= \sum_{m=-\infty}^{\infty} \left((\tau_{\theta\theta})_{00}^{(f)} F_m J_m(k_f r) + (\tau_{\theta\theta})_{10}^{(f)} F_m \frac{J_m'(k_f r)}{r} + (\tau_{\theta\theta})_{02}^{(f)} m^2 F_m \frac{J_m(k_f r)}{r^2} \right. \\ &\quad \left. + (\tau_{\theta\theta})_{00}^{(s)} S_m J_m(k_s r) + (\tau_{\theta\theta})_{10}^{(s)} S_m \frac{J_m'(k_s r)}{r} + (\tau_{\theta\theta})_{02}^{(s)} m^2 S_m \frac{J_m(k_s r)}{r^2} \right. \\ &\quad \left. + (\tau_{\theta\theta})_{01}^{(t)} m T_m \frac{J_m(k_t r)}{r^2} + (\tau_{\theta\theta})_{11}^{(t)} m T_m \frac{J_m'(k_t r)}{r} \right) e^{i(m\theta - \omega t)} \end{aligned} \quad (\text{C.115})$$

$$\begin{aligned} \tau_{r\theta} &= \sum_{m=-\infty}^{\infty} \left((\tau_{r\theta})_{01}^{(f)} m F_m \frac{J_m(k_f r)}{r^2} + (\tau_{r\theta})_{01}^{(s)} m S_m \frac{J_m(k_s r)}{r^2} + (\tau_{r\theta})_{10}^{(t)} T_m \frac{J_m'(k_t r)}{r} \right. \\ &\quad \left. + (\tau_{r\theta})_{11}^{(f)} m F_m \frac{J_m'(k_f r)}{r} + (\tau_{r\theta})_{11}^{(s)} m S_m \frac{J_m'(k_s r)}{r} + (\tau_{r\theta})_{02}^{(t)} m^2 T_m \frac{J_m(k_t r)}{r^2} + (\tau_{r\theta})_{20}^{(t)} T_m J_m''(k_t r) \right) e^{i(m\theta - \omega t)}. \end{aligned} \quad (\text{C.116})$$

The coefficients in the expansions for τ_{rr} are

$$\begin{aligned}
(\tau_{rr})_{00}^{(f)} &= \left(-A + Q \left(\frac{k_f^2 A_f}{\omega \sqrt{\rho'_{22}}} + \frac{\rho'_{12}}{\rho'_{22}} \right) + \omega \sqrt{\rho'_{22} A_f} \right) k_f^2 \\
(\tau_{rr})_{20}^{(f)} &= 2Nk_f^2 \\
(\tau_{rr})_{00}^{(s)} &= \left(-A + Q \left(\frac{k_s^2 A_s}{\omega \sqrt{\rho'_{22}}} + \frac{\rho'_{12}}{\rho'_{22}} \right) + \omega \sqrt{\rho'_{22} A_s} \right) k_s^2 \\
(\tau_{rr})_{20}^{(s)} &= 2Nk_s^2 \\
(\tau_{rr})_{01}^{(t)} &= -2iN \\
(\tau_{rr})_{11}^{(t)} &= 2iNk_t,
\end{aligned} \tag{C.117}$$

for $\tau_{\theta\theta}$ they are

$$\begin{aligned}
(\tau_{\theta\theta})_{00}^{(f)} &= \left(-A + Q \left(\frac{k_f^2 A_f}{\omega \sqrt{\rho'_{22}}} + \frac{\rho'_{12}}{\rho'_{22}} \right) + \omega \sqrt{\rho'_{22} A_f} \right) k_f^2 \\
(\tau_{\theta\theta})_{10}^{(f)} &= 2Nk_f \\
(\tau_{\theta\theta})_{02}^{(f)} &= -2N \\
(\tau_{\theta\theta})_{00}^{(s)} &= \left(-A + Q \left(\frac{k_s^2 A_s}{\omega \sqrt{\rho'_{22}}} + \frac{\rho'_{12}}{\rho'_{22}} \right) + \omega \sqrt{\rho'_{22} A_s} \right) k_s^2 \\
(\tau_{\theta\theta})_{10}^{(s)} &= 2Nk_s \\
(\tau_{\theta\theta})_{02}^{(s)} &= -2N \\
(\tau_{\theta\theta})_{01}^{(t)} &= 2iN \\
(\tau_{\theta\theta})_{11}^{(t)} &= -2iNk_t,
\end{aligned} \tag{C.118}$$

and for $\tau_{r\theta}$ they are

$$\begin{aligned}
(\tau_{r\theta})_{01}^{(f)} &= -2iN \\
(\tau_{r\theta})_{01}^{(s)} &= -2iN \\
(\tau_{r\theta})_{10}^{(t)} &= Nk_t \\
(\tau_{r\theta})_{11}^{(f)} &= 2iNk_f \\
(\tau_{r\theta})_{11}^{(s)} &= 2iNk_s \\
(\tau_{r\theta})_{02}^{(t)} &= -N \\
(\tau_{r\theta})_{20}^{(t)} &= -Nk_t^2.
\end{aligned} \tag{C.119}$$

The solid velocity $\dot{\mathbf{u}}$ is related to the spatial dependence part $\mathbf{u}^{(f)}$ of the time-harmonic solution by $\dot{\mathbf{u}} = -i\omega e^{-i\omega t} \mathbf{u}^{(f)}$. In terms of the Bessel function series expansion, using (C.70),

this gives

$$\dot{u}_r = \sum_{m=-\infty}^{\infty} \left((v_r)_{10}^{(f)} F_m J'_m(k_f r) + (v_r)_{10}^{(s)} S_m J'_m(k_s r) + (v_r)_{01}^{(t)} m T_m \frac{J_m(k_t r)}{r} \right) e^{i(m\theta - \omega t)} \quad (\text{C.120})$$

$$\dot{u}_\theta = \sum_{m=-\infty}^{\infty} \left((v_\theta)_{01}^{(f)} m F_m \frac{J_m(k_f r)}{r} + (v_\theta)_{01}^{(s)} m S_m \frac{J_m(k_s r)}{r} + (v_\theta)_{10}^{(t)} T_m J'_m(k_t r) \right) e^{i(m\theta - \omega t)} \quad (\text{C.121})$$

where the new coefficients are

$$(v_r)_{10}^{(f)} = -i\omega k_f, \quad (v_r)_{10}^{(s)} = -i\omega k_s, \quad (v_r)_{01}^{(t)} = \omega \quad (\text{C.122})$$

$$(v_\theta)_{01}^{(f)} = \omega, \quad (v_\theta)_{01}^{(s)} = \omega, \quad (v_\theta)_{10}^{(t)} = i\omega k_t. \quad (\text{C.123})$$

Finally, the relative fluid flow rate \mathbf{q} can be obtained as a function of position and time from $\mathbf{u}^{(f)}$ and $\mathbf{v}^{(f)}$ as $\mathbf{q} = -i\omega\phi e^{-i\omega t}(\mathbf{v}^{(f)} - \mathbf{u}^{(f)})$. Using (C.70) and (C.73) gives series expansions similar to those for the solid velocity,

$$q_r = \sum_{m=-\infty}^{\infty} \left((q_r)_{10}^{(f)} F_m J'_m(k_f r) + (q_r)_{10}^{(s)} S_m J'_m(k_s r) + (q_r)_{01}^{(t)} m T_m \frac{J_m(k_t r)}{r} \right) e^{i(m\theta - \omega t)} \quad (\text{C.124})$$

$$q_\theta = \sum_{m=-\infty}^{\infty} \left((q_\theta)_{01}^{(f)} m F_m \frac{J_m(k_f r)}{r} + (q_\theta)_{01}^{(s)} m S_m \frac{J_m(k_s r)}{r} + (q_\theta)_{10}^{(t)} T_m J'_m(k_t r) \right) e^{i(m\theta - \omega t)} \quad (\text{C.125})$$

where the radial fluid flow rate coefficients are

$$\begin{aligned} (q_r)_{10}^{(f)} &= ik_f \phi \left(\frac{k_f^2 A_f}{\sqrt{\rho'_{22}}} + \frac{\omega \rho'_{12}}{\rho'_{22}} + \omega \right) \\ (q_r)_{10}^{(s)} &= ik_s \phi \left(\frac{k_s^2 A_s}{\sqrt{\rho'_{22}}} + \frac{\omega \rho'_{12}}{\rho'_{22}} + \omega \right) \\ (q_r)_{01}^{(t)} &= -\omega \phi \left(1 + \frac{\rho'_{12}}{\rho'_{22}} \right) \end{aligned} \quad (\text{C.126})$$

and the azimuthal fluid flow rate coefficients are

$$\begin{aligned} (q_\theta)_{01}^{(f)} &= -\phi \left(\frac{k_f^2 A_f}{\sqrt{\rho'_{22}}} + \frac{\omega \rho'_{12}}{\rho'_{22}} + \omega \right) \\ (q_\theta)_{01}^{(s)} &= -\phi \left(\frac{k_s^2 A_s}{\sqrt{\rho'_{22}}} + \frac{\omega \rho'_{12}}{\rho'_{22}} + \omega \right) \\ (q_\theta)_{10}^{(t)} &= -i\omega k_t \phi \left(1 + \frac{\rho'_{12}}{\rho'_{22}} \right). \end{aligned} \quad (\text{C.127})$$

Finally, while the scattered component $p_w^{(r)}$ in the fluid is given directly by the expansion (C.100), the scattered component of fluid velocity may as well be explicitly written out in the time domain:

$$v_{wr}^{(r)} = \sum_{m=-\infty}^{\infty} \left((v_{wr}^{(r)})_{10} P_m H'_m(kr) \right) e^{i(m\theta - \omega t)} \quad (\text{C.128})$$

$$v_{w\theta}^{(r)} = \sum_{m=-\infty}^{\infty} \left((v_{w\theta}^{(r)})_{01} m P_m \frac{H_m(kr)}{r} \right) e^{i(m\theta - \omega t)} \quad (\text{C.129})$$

where

$$(v_{wr}^{(r)})_{10} = -\frac{ik}{\omega\rho_0} \quad (\text{C.130})$$

$$(v_{w\theta}^{(r)})_{01} = \frac{1}{\omega\rho_0} \quad (\text{C.131})$$

C.6.2 Procedure for evaluating the solution

With series representations for the variables of the first-order system available, a procedure can finally be stated for evaluating the solution of this scattering problem at any point (r, θ) and time t :

1. Choose an index L at which to truncate the series; the sums in (C.100) through (C.104) will range from $m = -L$ to L , rather than from $m = -\infty$ to ∞ .
2. For each m from $-L$ to L , solve (C.109) to obtain the series coefficients P_m , F_m , S_m , and T_m .
3. Use these series coefficients to evaluate the series representations of the first-order system variables from formulae (C.112), (C.114), (C.115), (C.116), (C.120), (C.121), (C.124), (C.125), (C.100), (C.128), and (C.129).

Note that the first two steps do not depend on where the solution is being evaluated; in particular, the series coefficients can be computed once per problem, then stored for later use.

C.7 Behavior of the series solution

The coefficients of this series solution have some properties that can make the solution easier to use. The series' convergence have also yet to be proven, and will be demonstrated here.

C.7.1 Negative indices

First, examine the behavior of the matrix \mathbf{M}_m and vector \mathbf{r}_m from equations (C.110) and (C.111) under the transformation $m \rightarrow -m$. Since m is an integer, the Bessel and Hankel functions transform as $J_{-m} = (-1)^m J_m$ and $H_{-m} = (-1)^m H_m$. The presence of a factor of

m in some terms of \mathbf{M}_m complicates matters, but the coefficient matrix can eventually be simplified to

$$\mathbf{M}_{-m} = (-1)^m \begin{pmatrix} (\mathbf{M}_m)_{11} & (\mathbf{M}_m)_{12} & (\mathbf{M}_m)_{13} & -(\mathbf{M}_m)_{14} \\ (\mathbf{M}_m)_{21} & (\mathbf{M}_m)_{22} & (\mathbf{M}_m)_{23} & -(\mathbf{M}_m)_{24} \\ (\mathbf{M}_m)_{31} & (\mathbf{M}_m)_{32} & (\mathbf{M}_m)_{33} & -(\mathbf{M}_m)_{34} \\ -(\mathbf{M}_m)_{41} & -(\mathbf{M}_m)_{42} & -(\mathbf{M}_m)_{43} & (\mathbf{M}_m)_{44} \end{pmatrix}. \quad (\text{C.132})$$

(Note that the bottom-left element is in fact zero, and is only written as $-(\mathbf{M}_m)_{41}$ for clarity.) This can be written more succinctly as

$$\mathbf{M}_{-m} = (-1)^m \mathbf{N} \mathbf{M}_m \mathbf{N}, \quad (\text{C.133})$$

where

$$\mathbf{N} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & -1 \end{pmatrix}, \quad (\text{C.134})$$

and conveniently $\mathbf{N}^{-1} = \mathbf{N}^T = \mathbf{N}$. Similarly, the right-hand side vector becomes

$$\mathbf{r}_{-m} = p_0 i^{-m} \begin{pmatrix} k J'_{-m}(ka) / (\omega^2 \rho_0) \\ -\eta_d J_{-m}(ka) \\ -J_{-m}(ka) \\ 0 \end{pmatrix} = (-1)^m i^{-2m} \mathbf{r}_m = \mathbf{r}_m. \quad (\text{C.135})$$

Thus the linear system for the coefficients \mathbf{c}_{-m} , $\mathbf{M}_{-m} \mathbf{c}_{-m} = \mathbf{r}_{-m}$, becomes

$$(-1)^m \mathbf{N} \mathbf{M}_m \mathbf{N} \mathbf{c}_{-m} = \mathbf{r}_m, \quad (\text{C.136})$$

whose solution is

$$\mathbf{c}_{-m} = (-1)^m \mathbf{N} (\mathbf{M}_m)^{-1} \mathbf{N} \mathbf{r}_m. \quad (\text{C.137})$$

Since the fourth component of \mathbf{r}_m is always zero, $\mathbf{N} \mathbf{r}_m = \mathbf{r}_m$, which gives

$$\begin{aligned} P_{-m} &= (-1)^m P_m \\ F_{-m} &= (-1)^m F_m \\ S_{-m} &= (-1)^m S_m \\ T_{-m} &= (-1)^{m+1} T_m. \end{aligned} \quad (\text{C.138})$$

This immediately also gives $T_0 = 0$.

With this result, and the assumption (yet to be proven!) of absolute convergence, the doubly-infinite series (C.100), (C.101), (C.102), and (C.103) can be converged to singly-

infinite ones,

$$p_w^{(r)} = P_0 H_0(kr) + 2 \sum_{m=1}^{\infty} P_{2m-1} H_{2m-1}(kr) \cos((2m-1)\theta) + 2 \sum_{m=1}^{\infty} P_{2m} H_{2m}(kr) \sin(2m\theta) \quad (\text{C.139})$$

$$\phi_f = F_0 J_0(k_f r) + 2 \sum_{m=1}^{\infty} F_{2m-1} J_{2m-1}(k_f r) \cos((2m-1)\theta) + 2 \sum_{m=1}^{\infty} F_{2m} J_{2m}(k_f r) \sin(2m\theta) \quad (\text{C.140})$$

$$\phi_s = S_0 J_0(k_s r) + 2 \sum_{m=1}^{\infty} S_{2m-1} J_{2m-1}(k_s r) \cos((2m-1)\theta) + 2 \sum_{m=1}^{\infty} S_{2m} J_{2m}(k_s r) \sin(2m\theta) \quad (\text{C.141})$$

$$\psi = 2 \sum_{m=1}^{\infty} T_{2m-1} J_{2m-1}(k_t r) \sin((2m-1)\theta) + 2 \sum_{m=1}^{\infty} T_{2m} J_{2m}(k_t r) \cos(2m\theta). \quad (\text{C.142})$$

Similar results could be obtained for the series representations of the variables of the first-order system.

C.7.2 Convergence

Let us now examine the behavior of the terms of the four series for large m and attempt to determine conditions under which they converge. Note that because of (C.138) it suffices to look at large positive m .

First, note the asymptotic behavior of the Bessel and Hankel functions for fixed argument as the order of the functions increases:

$$\left. \begin{aligned} J_m(z) &\sim \frac{1}{m!} \left(\frac{z}{2}\right)^m \\ H_m(z) &\sim -\frac{i(m-1)!}{\pi} \left(\frac{2}{z}\right)^m \end{aligned} \right\} \text{as } m \rightarrow \infty. \quad (\text{C.143})$$

From this, asymptotically as $m \rightarrow \infty$ the elements of the coefficient matrix \mathbf{M}_m behave

as

$$\begin{aligned}
(\mathbf{M}_m)_{11} &\sim -\frac{im!2^m}{\omega^2\rho_0\pi k^m a^{m+1}} \\
(\mathbf{M}_m)_{12} &\sim \frac{\delta_{1f}k_f^m a^{m-1}}{2^m(m-1)!} \\
(\mathbf{M}_m)_{13} &\sim \frac{\delta_{1s}k_s^m a^{m-1}}{2^m(m-1)!} \\
(\mathbf{M}_m)_{14} &\sim \frac{i\delta_{1t}k_t^m a^m}{2^m(m-1)!} \\
(\mathbf{M}_m)_{21} &\sim \frac{i\eta_d m!2^m}{\pi k^m a^m} \\
(\mathbf{M}_m)_{22} &\sim \frac{\delta_{2f}k_f^m a^{m-1}}{2^m(m-1)!} + \frac{\gamma_{2f}k_f^m a^m}{2^m m!} \\
(\mathbf{M}_m)_{23} &\sim \frac{\delta_{2s}k_s^m a^{m-1}}{2^m(m-1)!} + \frac{\gamma_{2s}k_s^m a^m}{2^m m!} \\
(\mathbf{M}_m)_{24} &\sim \frac{i\delta_{2t}k_t^m a^m}{2^m(m-1)!} \\
(\mathbf{M}_m)_{31} &\sim \frac{im!2^m}{\pi k^m a^m} \\
(\mathbf{M}_m)_{32} &\sim \frac{\epsilon_{3f}k_f^m a^{m-2}}{2^m(m-2)!} + \frac{\gamma_{3f}k_f^m a^m}{2^m m!} \\
(\mathbf{M}_m)_{33} &\sim \frac{\epsilon_{3s}k_s^m a^{m-2}}{2^m(m-2)!} + \frac{\gamma_{3s}k_s^m a^m}{2^m m!} \\
(\mathbf{M}_m)_{34} &\sim \frac{i\delta_{3t}k_t^m a^m}{2^m(m-1)!} + \frac{im\epsilon_{3t}k_t^m a^{m-1}}{2^m(m-1)!} \\
(\mathbf{M}_m)_{41} &= 0 \\
(\mathbf{M}_m)_{42} &\sim \frac{i\delta_{4f}k_f^m a^m}{2^m(m-1)!} + \frac{im\epsilon_{4f}k_f^m a^{m-1}}{2^m(m-1)!} \\
(\mathbf{M}_m)_{43} &\sim \frac{i\delta_{4s}k_s^m a^m}{2^m(m-1)!} + \frac{im\epsilon_{4s}k_s^m a^{m-1}}{2^m(m-1)!} \\
(\mathbf{M}_m)_{44} &\sim \frac{\delta_{4t}k_t^m a^{m-1}}{2^m(m-1)!} + \frac{\epsilon_{4tr}k_t^m a^{m-2}}{2^m(m-2)!} - \frac{\epsilon_{4t\theta}m k_t^m a^m}{2^m(m-1)!}.
\end{aligned} \tag{C.144}$$

Some common factors from the columns of this matrix can be extracted into a diagonal matrix multiplied on the right, obtaining

$$\mathbf{M}_m \sim \mathbf{M}'_m \mathbf{D}_r, \tag{C.145}$$

where

$$\mathbf{D}_r = \begin{pmatrix} \frac{im!2^m}{\pi k^m a^m} & 0 & 0 & 0 \\ 0 & \frac{k_f^m a^{m-2}}{2^m(m-1)!} & 0 & 0 \\ 0 & 0 & \frac{k_s^m a^{m-2}}{2^m(m-1)!} & 0 \\ 0 & 0 & 0 & \frac{k_t^m a^{m-2}}{2^m(m-1)!} \end{pmatrix} \quad (\text{C.146})$$

and, retaining only the $\mathcal{O}(m)$ and $\mathcal{O}(1)$ terms,

$$\mathbf{M}'_m = \begin{pmatrix} -\frac{1}{\omega^2 \rho_0 a} & \delta_{1f} a & \delta_{1s} a & i\delta_{1t} a^2 \\ \eta d & \delta_{2f} a & \delta_{2s} a & i\delta_{2t} a^2 \\ 1 & \epsilon_{3f} m - \epsilon_{3f} & \epsilon_{3s} m - \epsilon_{3s} & i\delta_{3t} a^2 + im\epsilon_{3t} a \\ 0 & i\delta_{4f} a^2 + im\epsilon_{4f} a & i\delta_{4s} a^2 + im\epsilon_{4s} a & \delta_{4t} a - \epsilon_{4tr} + \epsilon_{4tr} m - \epsilon_{4t\theta} a^2 m \end{pmatrix}. \quad (\text{C.147})$$

It is useful to further explicitly separate out the m dependence in \mathbf{M}'_m , defining $\mathbf{M}'_m =: \mathbf{M}_{m0} + m\mathbf{M}_{m1}$ where

$$\mathbf{M}_{m0} = \begin{pmatrix} -\frac{1}{\omega^2 \rho_0 a} & \delta_{1f} a & \delta_{1s} a & i\delta_{1t} a^2 \\ \eta d & \delta_{2f} a & \delta_{2s} a & i\delta_{2t} a^2 \\ 1 & -\epsilon_{3f} & -\epsilon_{3s} & i\delta_{3t} a^2 \\ 0 & i\delta_{4f} a^2 & i\delta_{4s} a^2 & \delta_{4t} a - \epsilon_{4tr} \end{pmatrix} \quad (\text{C.148})$$

$$\mathbf{M}_{m1} = \begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 \\ 0 & \epsilon_{3f} & \epsilon_{3s} & i\epsilon_{3t} a \\ 0 & i\epsilon_{4f} a & i\epsilon_{4s} a & \epsilon_{4tr} - \epsilon_{4t\theta} a^2 \end{pmatrix}. \quad (\text{C.149})$$

Note that it was necessary to retain the $\mathcal{O}(1)$ part of \mathbf{M}'_m as well as the $\mathcal{O}(m)$ part because the inverse of \mathbf{M}'_m is needed, and the $\mathcal{O}(m)$ part is singular. Retaining both these classes of terms is still consistent with the asymptotic expression (C.143) because the next most significant term is $\mathcal{O}(m^{-2})$ times smaller.

With these results, an asymptotic estimate for \mathbf{M}_m^{-1} can now be derived. First, decompose \mathbf{M}_{m1} explicitly as a rank-2 matrix, $\mathbf{M}_{m1} = \mathbf{U}\mathbf{V}^T$, where

$$\mathbf{U} = \begin{pmatrix} 0 & 0 \\ 0 & 0 \\ 1 & 0 \\ 0 & 1 \end{pmatrix}, \quad \mathbf{V}^T = \begin{pmatrix} 0 & \epsilon_{3f} & \epsilon_{3s} & i\epsilon_{3t} a \\ 0 & i\epsilon_{4f} a & i\epsilon_{4s} a & \epsilon_{4tr} - \epsilon_{4t\theta} a^2 \end{pmatrix}. \quad (\text{C.150})$$

Writing $\mathbf{M}'_m = \mathbf{M}_{m0} + \mathbf{U}(m\mathbf{I})\mathbf{V}^T$, and presuming that \mathbf{M}_{m0} is nonsingular, the Sherman-Morrison-Woodbury formula gives

$$(\mathbf{M}'_m)^{-1} = \mathbf{M}_{m0}^{-1} - \mathbf{M}_{m0}^{-1}\mathbf{U}(m^{-1}\mathbf{I} + \mathbf{V}^T\mathbf{M}_{m0}^{-1}\mathbf{U})^{-1}\mathbf{V}^T\mathbf{M}_{m0}^{-1}, \quad (\text{C.151})$$

or, asymptotically as $m \rightarrow \infty$,

$$(\mathbf{M}'_m)^{-1} \sim \mathbf{M}_{m0}^{-1} - \mathbf{M}_{m0}^{-1}\mathbf{U}(\mathbf{V}^T\mathbf{M}_{m0}^{-1}\mathbf{U})^{-1}\mathbf{V}^T\mathbf{M}_{m0}^{-1} =: \mathbf{M}_\infty^{-1}. \quad (\text{C.152})$$

Thus as $m \rightarrow \infty$, the inverse of the coefficient matrix \mathbf{M}_m can be asymptotically described as

$$\mathbf{M}_m^{-1} \sim \mathbf{D}_r^{-1} \mathbf{M}_\infty^{-1}. \quad (\text{C.153})$$

The significance of this is that \mathbf{M}_∞^{-1} is a constant matrix with respect to m , and \mathbf{D}_r^{-1} applies a simple scaling factor to each coefficient.

Turning to the right-hand side vector \mathbf{r}_m , using the asymptotic expression for J_m from (C.143), we get

$$\mathbf{r}_m \sim p_0 i^m \begin{pmatrix} \frac{k^m a^{m-1}}{\omega^2 \rho_0 2^m (m-1)!} \\ -\frac{\eta_d k^m a^m}{2^m m!} \\ -\frac{k^m a^m}{2^m m!} \\ 0 \end{pmatrix} = p_0 \left(\frac{ika}{2} \right)^m \frac{1}{m!} \begin{pmatrix} \frac{m}{\omega^2 \rho_0 a} \\ -\eta_d \\ -1 \\ 0 \end{pmatrix}. \quad (\text{C.154})$$

Explicitly extracting out the m dependence, let

$$\mathbf{r}_{m0} := (0 \quad -\eta_d \quad -1 \quad 0)^T, \quad \mathbf{r}_{m1} := \left(\frac{1}{\omega^2 \rho_0 a} \quad 0 \quad 0 \quad 0 \right)^T, \quad (\text{C.155})$$

which gives

$$\mathbf{r}_m \sim p_0 \left(\frac{ika}{2} \right)^m \frac{1}{m!} (\mathbf{r}_{m0} + m \mathbf{r}_{m1}). \quad (\text{C.156})$$

Combining the asymptotic expressions for \mathbf{M}_m and \mathbf{r}_m yields

$$\mathbf{c}_m \sim \mathbf{D}_r^{-1} \mathbf{M}_\infty^{-1} \left(p_0 \left(\frac{ika}{2} \right)^m \frac{1}{m!} \right) (\mathbf{r}_{m0} + m \mathbf{r}_{m1}). \quad (\text{C.157})$$

Absorbing the scalar factors originating from the \mathbf{r}_{m1} term into \mathbf{D}_r^{-1} and dropping the \mathbf{r}_{m0} term results in

$$\mathbf{c}_m \sim p_0 \begin{pmatrix} i^{m-1} \pi \left(\frac{ka}{2} \right)^{2m} \frac{1}{m!(m-1)!} & 0 & 0 & 0 \\ 0 & a^2 \left(\frac{ik}{k_f} \right)^m & 0 & 0 \\ 0 & 0 & a^2 \left(\frac{ik}{k_s} \right)^m & 0 \\ 0 & 0 & 0 & a^2 \left(\frac{ik}{k_s} \right)^m \end{pmatrix} \mathbf{M}_\infty^{-1} \mathbf{r}_{m1}. \quad (\text{C.158})$$

Now all the asymptotic m dependence has been absorbed into the leading diagonal matrix. There are thus constants P_∞ , F_∞ , S_∞ , and T_∞ such that the coefficients can be essentially written as

$$\begin{aligned} P_m &\sim P_\infty i^{m-1} \left(\frac{ka}{2} \right)^{2m} \frac{1}{m!(m-1)!} \\ F_m &\sim F_\infty \left(\frac{ik}{k_f} \right)^m \\ S_m &\sim S_\infty \left(\frac{ik}{k_s} \right)^m \\ T_m &\sim T_\infty \left(\frac{ik}{k_t} \right)^m. \end{aligned} \quad (\text{C.159})$$

This constitutes something of an abuse of the asymptotic \sim notation, because the possibility cannot be ruled out that the constants P_∞ etc. may be zero; the notation that $f(m) \sim g(m)$ as $m \rightarrow \infty$ means that $\lim_{m \rightarrow \infty} f(m)/g(m) = 1$, which is nonsensical if $g(m) \equiv 0$. If this is the case, though, the corresponding coefficient grows more slowly with m than the given estimate, so the convergence results below are not jeopardized.

Finally, from these asymptotic expressions for the series coefficients, the convergence of the series can now be examined. Term m of the series for $p_w^{(r)}$ is asymptotically

$$P_m H_m(kr) e^{im\theta} \sim -\frac{P_\infty e^{im\theta}}{\pi} \left(\frac{ika^2}{2r}\right)^m \frac{1}{m!}. \quad (\text{C.160})$$

By the Ratio Test, then, the series (C.100) for $p_w^{(r)}$ is absolutely convergent for all $r > 0$, and uniformly convergent for $r \geq a$. Similarly, for the other series,

$$\{F, S, T\}_m J_m(k_{\{f,s,t\}} r) e^{im\theta} \sim \{F, S, T\}_\infty e^{im\theta} \left(\frac{ikr}{2}\right)^m \frac{1}{m!}, \quad (\text{C.161})$$

which again gives absolute convergence for all finite r by the Ratio Test, and uniform convergence for $r \leq a$. In fact, expressions (C.160) and (C.161) imply even more than just absolute convergence of the series – by differentiating with respect to r or θ , they also give asymptotic results that imply the absolute convergence of all the derivatives of the potentials $p_w^{(r)}$, ϕ_f , ϕ_s , and ψ , and uniform convergence on the domains where these potentials are defined, meaning that the series expressions formally derived for the variables of the first-order system are in fact absolutely and uniformly convergent.

Appendix D

GALLERY OF ENERGY DENSITY PLOTS FOR CYLINDRICAL SCATTERER CASES

This appendix contains plots of the energy densities at the initial time for all 18 cases selected for convergence investigation in Section 7.1, computed using the analytical solution formula developed in Appendix C. The plots here also include locations and values of maxima of the energy density within the scatterer, which were omitted from the main text due to the difficulty of including them legibly in a small plot. The plotting is the same as in Figure 7.2: the plotted quantity is the energy density, normalized by the peak energy density of the incident wave, and contours are placed at powers of 2, with dashed contours indicating negative powers and solid indicating positive powers. The blue \times symbols indicate the locations of maxima, with the adjacent labels giving the value at those maxima. Since the solutions are symmetric top-to-bottom, only the top or bottom maximum of a pair is labeled, not both.




Figure D.1: Cylindrical scatterer case 1: Inviscid, $\eta_d = 0$, 13.25 kHz




Figure D.2: Cylindrical scatterer case 2: Inviscid, $\eta_d = 0$, 22.25 kHz

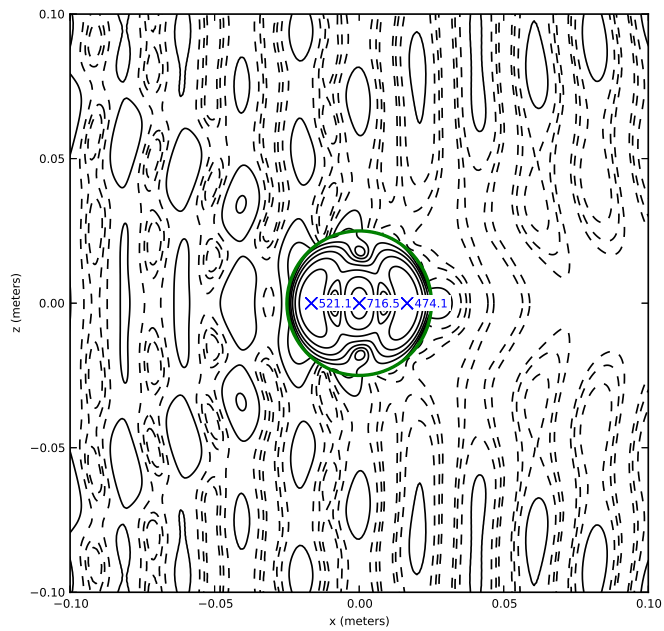


Figure D.3: Cylindrical scatterer case 3: Inviscid, $\eta_d = 0$, 38.20 kHz




Figure D.4: Cylindrical scatterer case 4: Inviscid, $\eta_d = 0.5$, 17.25 kHz




Figure D.5: Cylindrical scatterer case 5: Inviscid, $\eta_d = 0.5$, 25.02 kHz




Figure D.6: Cylindrical scatterer case 6: Inviscid, $\eta_d = 0.5$, 39.03 kHz




Figure D.7: Cylindrical scatterer case 7: Inviscid, $\eta_d = 1$, 17.30 kHz




Figure D.8: Cylindrical scatterer case 8: Inviscid, $\eta_d = 1$, 25.09 kHz




Figure D.9: Cylindrical scatterer case 9: Inviscid, $\eta_d = 1$, 39.04 kHz




Figure D.10: Cylindrical scatterer case 10: Viscous, $\eta_d = 0$, 15.70 kHz




Figure D.11: Cylindrical scatterer case 11: Viscous, $\eta_d = 0$, 24.55 kHz




Figure D.12: Cylindrical scatterer case 12: Viscous, $\eta_d = 0$, 38.39 kHz




Figure D.13: Cylindrical scatterer case 13: Viscous, $\eta_d = 0.5$, 15.80 kHz




Figure D.14: Cylindrical scatterer case 14: Viscous, $\eta_d = 0.5$, 24.54 kHz




Figure D.15: Cylindrical scatterer case 15: Viscous, $\eta_d = 0.5$, 38.35 kHz




Figure D.16: Cylindrical scatterer case 16: Viscous, $\eta_d = 1$, 15.80 kHz




Figure D.17: Cylindrical scatterer case 17: Viscous, $\eta_d = 1$, 24.54 kHz




Figure D.18: Cylindrical scatterer case 18: Viscous, $\eta_d = 1$, 38.35 kHz

VITA

Grady Lemoine is a man of the Pacific Northwest by birth and inclination. He obtained a Bachelor's degree from the Department of Aeronautics and Astronautics at the University of Washington in 2002, and a Master's degree from the same department in 2004. Following this, he worked at Boeing Commercial Aircraft for several years, before returning to the University of Washington to seek higher education, this time from the Department of Applied Mathematics. He earned the degrees of Master of Science and Doctor of Philosophy in Applied Mathematics in 2008 and 2013, respectively.