

An Extensible Computational Cochlear Modelling Framework

Michael James Rapson

A thesis submitted to the Department of Electrical Engineering,
University of Cape Town, in fulfilment of the requirements
for the degree of Doctor of Philosophy.

Cape Town, February 2012

Declaration

I declare that this thesis is my own work in substance. All assistance that I have received is gratefully acknowledged, and detailed on page iv. It is being submitted for the degree of Doctor of Philosophy in the University of Cape Town. It has not been submitted before for any degree or examination in any other university.

Signature of Author

Cape Town

10 February 2012

Abstract

Cochlear models are used by many people for many purposes. Each purpose has different accuracy and complexity requirements, so no single model will be suitable for every purpose. A wide range of cochlear models has already been proposed, and the diversity in the field is only increasing. The proposed framework allows a large subset of models to be implemented and compared, facilitating modelling efforts. It uses the finite element method and state space techniques in an object oriented library, and is suitable for time domain cochlear models. In the face of the diversity noted, the physiology of the cochlea itself limits the scope of the possible models to consider. A general modelling framework is described in broad strokes, narrowing to provide full details of its implementation and use for a smaller set of models. Various published models were accurately replicated in the framework, and it was used to verify a silicon cochlea model. Studies into the mathematics behind fluid-structure interaction in the cochlea show that the set of models considered is characteristic of cochlear models in general, ensuring the desired extensibility of the modelling framework. These studies also have various implications for the cochlear modelling community at large.

Acknowledgements

As I approach the end of this project, I have occasion to think about all the people have helped me to reach this point. It has been said that without counsel plans fail, but with many advisers they succeed. I can certainly agree.

Jonathan Tapson has been my supervisor in this project. I have great respect for his abilities as a researcher, supervisor, lecturer, and entrepreneur. This made him my natural choice to work with. I would like to thank him for taking me on as a student, and directing me towards cochlear modelling as a field that I was suited to. Furthermore, I would like to wish him all the best at his new post at the University of Western Sydney. I look forward to working together in the future.

The cochlear modelling community has been welcoming and generous with their ideas. I have had the opportunity to meet or correspond with: Jont Allen; Dick Lyon; Christopher Shera; Peter van Hengel; John Matthews; Hendrikus Duifhuis; Matthew Flax; Andrew Bell; Stephen Neely; Martin Braun; and Sripriya Ramamoorthy, to mention a few people by name. I appreciate their input to my project. An incident that stands out in my memory is that John Matthews, who transitioned to another field in the 1980s, offered to scan his 200 page thesis for me when I enquired about his work.

I have had the opportunity to collaborate with Tara Hamilton and André van Schaik in a more formal manner. In addition to email correspondence and meetings during short visits, I spent May and June 2010 working with them in Sydney. During this trip, the foundation for the work in chapter 8 was laid.

Closer to home, Daya Reddy, head of CERECAM, provided guidance when I was considering using the finite element method heavily in my work; Mike Inggs helped me to access the CSIR Centre for High Performance Computing; and Martin Braae, my undergraduate supervisor, continued to provide advice and encouragement throughout the project. I made extensive use of the CERECAM cluster. I found Graham Inggs, the administrator of the cluster, exceptionally helpful.

Victor Udoewa was my first contact in CERECAM. He introduced me to Daya Reddy as well as to the deal.II software library. I have learnt a lot about the finite element method and deal.II from the deal.II users group at UCT. I must thank Andrew McBride, Andrew McMahon, J-P Pelteret, and Beverley Grieshaber in particular. I must also extend my thanks to the deal.II authors: Wolfgang Bangerth; Ralf Hartmann; and Guido Kanschat. I have found deal.II to be an excellent library and

particularly appreciate the way that questions on the mailing list are answered promptly, often within an hour of the question being asked – I have made frequent use of this.

Technical advice and assistance alone would have been insufficient to complete this work. I have counted on my friends and family in a number of ways throughout this work. (This is not to imply that the people who I have already mentioned are not also friends in their own right!) I would like to mention them now.

J-P Kloppers, Andrew Purchase, Nkulu Mudonko, Mark Matthews, and Grant Smith have all provided invaluable guidance throughout my time in Cape Town. My flatmates over the years have become some of my most trusted friends. I would like to mention Ross Douglas, Peter Waller, Duane Harvey, and Fiona Holtz in particular. Warren Gatcke also manages to sneak into this group. Some of the people I studied with have also become great friends – Ross Allardice, David Karpul, Kim McMahon, Ashley Liddiard, and Neil Bradshaw must be mentioned.

Sally Marks and Mark Albertyn were my high school maths and physics teachers. Their classes were enjoyable, but they also laid firm foundations beyond the subject matter they presented.

My older brother Trevor has gone before me in every area of life, and I am glad to have followed in his footsteps. I have learnt a lot from him in many ways. Thank you for your example and advice.

I have been blessed with exceptional parents. Their lives both as individuals and as a couple has been, and will continue to be, an invaluable example to me. I am sure that I am only partially aware of how much I have learnt and picked up from them. I value my father's advice and guidance highly and am convinced that I learnt the mindset that I have needed for this work from him. Although I am proud to take after my father in many obvious ways, my mother has also had an essential influence on my life. She has consistently been caring and concerned about my progress through life, encouraging me and challenging me as needed. She is an example of strength, having "stuck it out" in Zimbabwe through its worst days.

Finally, I would like to acknowledge the financial assistance of the NRF (South Africa), the University of Cape Town Doctoral Research Scholarship, and the Harry Crossley Foundation towards my personal expenses during this project. Funding from the Max & Lillie Sonnenberg Scholarship and the Crasnow Postgraduate Scholarship allowed me to work with André and Tara in Sydney. I also received a DAAD Local Scholarship that I ultimately had to turn down due to restrictions on holding NRF and DAAD funding simultaneously. The offer was nevertheless appreciated.

Contents

Declaration	i
Abstract	iii
Acknowledgements	iv
Nomenclature	xvi
1 Introduction	1
1.1 Motivation for the Current Study	1
1.2 Parallel Tracks	2
1.3 Comments on Structure	4
1.4 Contribution	4
1.5 Related Publications	5
1.6 Recommended Reading	6
2 Models	8
2.1 What is a Model?	8
2.1.1 General Usage of the Term	8
2.1.2 Abstract Models and Implemented Models	11
2.1.3 Summary	13
2.2 One Organ, Many Motives for Studying It	14
2.2.1 Why Study the Cochlea?	14
2.2.2 Differing Requirements Each Use Places on the Models	18

2.2.3	Fields Relevant to Cochlear Modelling	20
2.2.4	Summary	21
2.3	Statement of Problem	21
2.3.1	A History of Cochlear Modelling Frameworks	22
2.3.2	An Object-Orientated Computational Modelling Framework	23
2.3.3	Power and Flexibility versus Safety and Efficiency	24
2.3.4	What Won't be Covered	25
2.3.5	Definitions Used in this Thesis	26
2.4	Summary	26
3	An Introduction to the Cochlea	28
3.1	Location in the Ear	28
3.1.1	The Outer Ear	30
3.1.2	The Middle Ear	30
3.1.3	The Auditory Pathway in the Brainstem	31
3.2	Anatomy of the Cochlea	33
3.2.1	The Scalae	33
3.2.2	The Cochlear Partition	34
3.2.3	Innervation	36
3.3	Physiology of the Cochlea	38
3.3.1	Vulnerability	38
3.3.2	Definition of Passive and Active Behaviour	39
3.3.3	Passive Behaviour	41
3.3.4	Active Behaviour	44
3.3.5	Hair Cell Physiology	50
3.4	Species Dependence	54
3.5	Ethics	56
3.6	Summary	58

4	Specification – Implications of Cochlear Physiology	59
4.1	Objects and Inheritance	59
4.2	The Cochlear Fluid	61
4.2.1	Incompressibility – Viergever’s Contribution	62
4.2.2	Linearity – Viergever’s Contribution	63
4.2.3	Viscosity – Viergever’s Contribution	63
4.2.4	Viergever’s Contribution – Implications	64
4.2.5	Viergever’s Contribution – Reassessment	65
4.2.6	Spatial Dimension	66
4.2.7	Fluid Geometry	68
4.2.8	Fluid-Structure Interaction and Mesh Moving	69
4.2.9	Conclusion	70
4.3	State Space and Time Integration	70
4.3.1	Definitions	70
4.3.2	Formal Limitations on these Methods	73
4.3.3	Strengths and Weaknesses	75
4.3.4	Alternatives	77
4.4	The Cochlear Partition	80
4.4.1	Passive Dynamics	83
4.4.2	Active Mechanisms	86
4.4.3	Physical Coupling between Segments	88
4.4.4	Innervation	91
4.4.5	Cochlear Partition Compilation	92
4.5	Input Boundaries	93
4.5.1	The Outer Ear, Middle Ear, and Stapes	93
4.5.2	Bone Conduction	95
4.6	Other Boundaries	95
4.6.1	The Round Window	95
4.6.2	The Helicotrema	97

4.6.3	The “Third” Window	98
4.7	Parallelisation	99
4.8	Summary	102
5	Fluid-Structure Interaction	105
5.1	The Monolithic Approach	105
5.1.1	Key Results	106
5.1.2	Implications for the Modelling Framework	110
5.2	More on Nonlinearity	112
5.2.1	The Nonlinear Fluid-Structure Feedback Loop	112
5.2.2	Linearisation of the State Update Equation	113
5.2.3	Dependence of \mathbf{A}_E , \mathbf{B}_E , and \mathbf{Y}_E on $\mathbf{p}(t)$ and $\mathbf{X}(t)$	115
5.2.4	Implications for the Modelling Framework	117
5.3	A Partitioned Approach?	117
5.3.1	Key Results	118
5.3.2	Implications for the Modelling Framework	120
5.4	Summary	121
6	Detailed Description of the Framework	123
6.1	Introduction	123
6.2	The Fluid Model	123
6.2.1	Sign Conventions	124
6.2.2	Segment Input/Output Conventions	124
6.2.3	Weakform Derivation	125
6.2.4	Function Space Details	128
6.3	Focusing on the Organ of Corti Complex	129
6.3.1	Fluid Inputs to Output Linearisation	130
6.3.2	Other Inputs to Output Linearisation	131
6.3.3	Data Outputs	133
6.3.4	Inner Hair Cell and Nerve Models	134

6.3.5	Locating Required Inputs and Outputs	134
6.3.6	Conventions Regarding Data Outputs and External Inputs	134
6.4	Other Boundary Conditions	135
6.5	Summary	135
7	A First Application	136
7.1	Introduction	136
7.1.1	Models with Emergent Travelling Wave Behaviour	136
7.1.2	Ways Technology Projects Can Fail	137
7.1.3	Common Criticisms of the Travelling Wave Hypothesis	140
7.2	Segment Coupling through the Cochlear Fluid	145
7.2.1	Classical Cochlear Partition Models	145
7.2.2	Isolating Segment Mass	147
7.2.3	Behaviour for Mass Limits	149
7.2.4	Discussion	151
7.3	Steps to Build the Model	152
7.3.1	Select a Cochlear Partition Model Type	153
7.3.2	Select Tuning Parameters	156
7.3.3	Specify a Fluid Region	159
7.3.4	Choose Boundary Conditions	161
7.3.5	Assemble the Pieces	162
7.4	Results	163
7.4.1	Steady State Behaviour	163
7.4.2	Transient Behaviour	164
7.4.3	Postprocessing and Visualisation	166
7.4.4	Results from a Two Scalae Model	171
7.5	Discussion	172
7.6	Summary	175

8 Hopf, Van der Pol, and Silicon Cochleae	176
8.1 Background	177
8.1.1 Van der Pol Oscillator Models	177
8.1.2 Hopf Bifurcation Models	178
8.1.3 Application to Silicon Cochleae	179
8.2 Implementation in the Proposed Framework	180
8.2.1 Select a Cochlear Partition Model Type	180
8.2.2 Select Tuning Parameters	182
8.2.3 Specify a Fluid Region	182
8.2.4 Choose Boundary Conditions	182
8.2.5 Assemble the Pieces	183
8.3 Comparisons Between the Silicon Cochlea and the Modelling Framework	183
8.3.1 Cochlear Fluid	183
8.3.2 Cochlear Partition Model	184
8.3.3 Four Coupled Segments	190
8.4 Operating Cochlear Models Beyond the Bifurcation Point	192
8.5 Summary	194
8.5.1 VLSI Implementation and Hopf Bifurcation Models	194
8.5.2 Van der Pol Oscillator Models	195
8.5.3 Support for Central Ideas in this Thesis	195
9 Conclusions	197
9.1 An Extensible Cochlear Modelling Framework	197
9.2 Models Implemented in the Framework To Date	198
9.3 Supporting Derivations and Results	199
A Modes of OAE Observation	201

B	Results from Parallel Processing Tests	203
B.1	Introduction	203
B.2	Method	204
B.3	Results	205
B.4	Discussion	207
C	Unification and Extension of Monolithic State Space and Iterative Cochlear Models	208
D	Investigations into Time Stepping Methods for Cochlear Models	227
E	Capturing the Low Frequency Cochlear Impedance in Time Domain Models: The Role of Viscosity	254
E.1	Introduction	254
E.2	Original Article Abstract	255
E.3	Slides Presented	255
E.4	Presentation Script	257
F	Ethical Assessment	261
	Bibliography	263

List of Figures

2.1	Mechanical cochlear model constructed by George Wilkinson.	9
2.2	A network model of the cochlea from 1951, and a silicon cochlear model from 2008.	10
2.3	A triangle and two methods to approximate it, which can serve as an analogy to the modelling process.	11
3.1	Diagram of the ear.	29
3.2	Diagram of the main ascending pathways of the brainstem.	31
3.3	Centrifugal pathways from the superior olivary complex to the cochlea.	32
3.4	Cross-section through a turn of the cochlea.	34
3.5	Cross-section through the organ of Corti.	35
3.6	Diagrams of the inner and outer hair cells.	35
3.7	Diagram showing the innervation of the organ of Corti.	37
3.8	Targets of the crossed and uncrossed olivocochlear bundle in the cochlea.	38
3.9	Diagram showing the effect of death on the basilar membrane response.	40
3.10	Georg von Békésy's figures showing the frequency response of the cochlear partition.	42
3.11	Donald Greenwood's original mapping of best place with respect to frequency. . . .	43
3.12	Diagram showing the nonlinear response of the cochlea to stapes input.	45
3.13	Comparison between the basilar membrane response and the neural tuning curve. . .	46
3.14	Diagram showing two-tone suppression in the basilar membrane response.	47
3.15	Diagram showing distortion products in the basilar membrane response.	49
3.16	First evidence of otoacoustic emissions published by David Kemp.	51
3.17	Motility in the outer hair cell.	53
3.18	Transducer currents in the outer hair cell.	55

3.19	The voltage response of the inner hair cell.	55
4.1	An inheritance diagram for the cochlear modelling framework.	60
4.2	Methods of implementing impedance in the time domain.	79
4.3	Simplified cross-section of the organ of Corti complex.	81
4.4	Diagram showing the arrangement of outer hair cells in a radial segment.	89
4.5	An inheritance diagram for the cochlear partition.	103
4.6	An inheritance diagram for the OtherBoundary class.	104
6.1	Diagram showing typical finite element cells at the boundary between organ of Corti complex and the fluid.	126
7.1	The effect of segment mass on fluid-structure coupling.	152
7.2	The envelope of the cochlear partition response for high and low segment mass conditions.	159
7.3	Steady state response for high and low segment mass.	163
7.4	The transient response of isolated segments.	164
7.5	The transient response for high segment mass.	165
7.6	The transient response for low segment mass.	165
7.7	The velocity of the cochlear partition at five time instants for a cochlear model with high segment mass.	167
7.8	The velocity of the cochlear partition at five time instants for a cochlear model with low segment mass.	168
7.9	Postprocessed data for a model with high segment mass.	169
7.10	Postprocessed data for a model with low segment mass.	170
7.11	Postprocessed data for a model with low segment mass showing vertical component of the intensity.	170
7.12	Steady state response for a model with two separate fluid scalae.	171
7.13	Postprocessed data from the model with two separate fluid scalae.	172
8.1	Cadence simulations of a single segment for a sine wave input.	185
8.2	Numerical simulations of a single segment for a sine wave input.	186

8.3	An example of the extreme distortion that Hopf bifurcation and Van der Pol oscillator models can allow for large inputs.	187
8.4	The effect of the gain parameter on the Cadence simulation and the framework implementation.	188
8.5	Results from the biological cochlea and a silicon cochlea that show linearisation in the response to high stimulus intensities.	189
8.6	Cadence simulation of four radial segments coupled by fluid.	190
8.7	Numerical simulation of four radial segments coupled by fluid.	191
8.8	Response from a Van der Pol oscillator model.	193
B.1	A comparison of fluid solver time for various processors and solvers.	206

Nomenclature

1D (etc.) One-dimensional (etc.)

Active system A system that contains energy sources. However, if the energy sources present only partially offset losses in the system, then it can share many properties of a passive system, and may be referred to as passive in a loose sense. The term “locally active” has been used to distinguish systems where the energy sources more than offset any losses present.

Basilar membrane The chief stiffness element in the cochlear partition, often used as a synonym for “cochlear partition” or “organ of Corti complex” in this thesis. See chapter 3 for its formal definition.

Best frequency The frequency that evokes the largest response at a particular location on the basilar membrane or cochlear partition.

Best place The location on the basilar membrane or cochlear partition that is most sensitive to a particular input frequency.

Bifurcation An abbreviation of Hopf bifurcation used in this thesis.

Cadence[®] A VLSI silicon chip design and simulation software package.

Characteristic frequency The frequency that evokes the largest response at a particular place on the cochlear partition, or in a particular nerve fibre.

CGS Centimetre-gram-second.

Cochlear partition The structure dividing the cochlea into two main channels along its length, often used as a synonym for “basilar membrane” or “organ of Corti complex” in this thesis. See chapter 3 for its formal definition.

Critical value The values of tuning parameters that place a system precisely at a Hopf bifurcation.

d.c. Direct current.

dB Decibel. A logarithmic scale based on ratios, defined as $10 \log_{10}(A/A_r)$ for quantities related to power, and $20 \log_{10}(A/A_r)$ for quantities related to amplitude. A_r is a reference value required by the scale.

deal.II An open-source C++ based finite element library (Differential Equations Analysis Library).

dyn Dyne, a measure of force in the CGS unit system. $1 \text{ dyn} = 1 \text{ g} \cdot \text{cm} \cdot \text{s}^{-2} = 10^{-5} \text{ N}$.

erg A unit of energy in the CGS system. $1 \text{ erg} = 1 \text{ g} \cdot \text{cm}^2 \cdot \text{s}^{-2} = 10^{-7} \text{ J}$.

flops Floating point operations per second.

Fluid-structure interaction The behaviour that occurs at the interface between a fluid and a deformable structure

Framework A specification of the basic components and structure required to support the implementation of a model. Computer code or otherwise that provides these things.

Hopf bifurcation Systems that change from a stable response to an oscillatory response for some set of tuning parameters contain a Hopf bifurcation. The set of parameters for which the transition occurs is called the critical value. If the transition is smooth, the system undergoes a *supercritical* Hopf Bifurcation. The transition is not smooth for *subcritical* bifurcations.

Implementation A specific realisation of an abstract model from which results can be obtained.

Macroscopic model A cochlear model that represents the whole cochlea but with low detail.

MEMS Micro-Electro-Mechanical Systems.

MI Monolithic iterative. See section 5.1.1.3.

Microscopic model A cochlear model that represents a small part of the cochlea (especially the cochlear partition) in detail.

MKS Metre-kilogram-second.

Model Mathematical construct representing a physical object or system.

Monolithic model A numerical modelling approach where all parts of the model contribute towards a single, large matrix equation.

MSS Monolithic state space. See section 5.1.1.3.

MSSV Monolithic state space variation. See section 5.1.1.3.

Organ of Corti complex A collective word for the various structures around the organ of Corti proper. Often used as a synonym for “basilar membrane” or “cochlear partition” in this thesis. See chapter 3 for its formal definition.

Partition An abbreviation of the term “cochlear partition”.

Partitioned model A numerical modelling approach where the model is implemented as separate parts that exchange data during the simulation.

Partitioning scheme The approach to setting up a partitioned model.

Passive system A system that does not contain or require any energy sources. However, if energy sources are present, but only partially offset losses in the system, then it can share many properties of a passive system, and may be referred to as passive in a loose sense.

Q Quality factor. See section 7.3.2.

Radial segment A short section of the basilar membrane, cochlear partition, or organ of Corti complex. These structures are modelled by dividing them into a number of radial segments. The anatomy of these structures encourages this approach, see chapters 3 and 4.

RAM Random access memory.

re Referred to. Used when quoting pressures using SPL units.

Real-time Simulations which take as long to complete as the simulated input duration are said to run in “real-time”.

Run-time The time required for a computer simulation to complete.

State “Any complete independent set of system model variables whose knowledge enables one to predict completely the future behaviour of the system” (MacFarlane 1970, chapter 6).

Segment An abbreviation of the term “radial segment”.

Silicon cochlea A cochlear model implemented using VLSI technology.

SPL Sound Pressure Level. A decibel scale for sound pressure: $20 \log_{10} (p_{\text{rms}}/p_{\text{ref}})$. It is generally referred to $20 \mu\text{Pa}$, although other values are sometimes used.

Travelling wave hypothesis The hypothesis that a travelling wave, which is visible on the basilar membrane, is primarily responsible for the response of the ear to sound.

VLSI Very Large Scale Integration – referring to silicon chips containing large numbers of transistors.

Chapter 1

Introduction

1.1 Motivation for the Current Study

The cochlea is the section of the ear that converts sound-evoked displacements into neural impulses to be transferred to the brain. It has been the subject of research since at least the 16th century (Duifhuis 2012). The pace of research accelerated from the year 1800, with various forms of cochlear models supporting the work. This thesis identifies the current and future users of cochlear models, and describes their requirements for the models, as well as how they will make use of them. Finally, and most importantly, this thesis studies the steps we can take today to facilitate the models of the future.

It seems reasonable to assume that one approach to cochlear modelling will ultimately win the floor, and all the alternatives will simply be viewed as stepping stones along the path. However, this scenario is unlikely given the diversity of people who will use cochlear models, and the roles the model will fill. Certainly, some models will be (and have been) shown to be inherently wrong. This is the nature of science. Many others will be outperformed by a few percent in terms of accuracy, making them unsuitable for some applications, but still ideal for others. This is the nature of engineering. Cochlear modelling takes place on the boundary between science, medicine, and engineering. It draws interest from all of these areas. We still have much to learn about the cochlea, and this is reflected in the wide range of models that have been, and still are being, considered. I will argue that a means to construct and compare different models within the same framework will facilitate future models. For this to be successful, the framework should allow the widest possible set of models to be compared fairly.

The need to be able to compare models directly and fairly was brought home to me by two separate things. Firstly, this project started with the aim of supporting work on Very Large Scale Integration (VLSI) silicon cochlea models by our collaborators André van Schaik and Tara Hamilton in Sydney. Could numerical simulations of cochlear models be used to suggest features to include in the

chips? A numerical cochlear modelling framework that could compare models seemed an obvious requirement. Chapter 8 shows how the framework was ultimately used in this regard.

Secondly, cochlear modelling researchers seem polarised into two camps along the lines of whether they agree with the travelling wave hypothesis or not. (Most do.) Their positions are understandable: the travelling wave hypothesis has won acceptance over numerous alternatives through the combination of extensive physiological results and modelling efforts showing its plausibility. Researchers cannot be expected to simply ignore or discard the progress made with these models. On the other hand, people who object to the travelling wave hypothesis typically point to some areas where travelling wave models fail to explain the physiological data. Having personally investigated this area and found an alternative model that they think explains the data better, they are frustrated if the model they propose is not taken seriously. I personally find many of the mechanisms they propose intriguing. In particular, I wonder whether there is not some middle ground, maybe a scenario where some of the mechanisms proposed as an alternative to the travelling wave are ultimately adopted into a hybrid model of some form. The current polarisation works against this possibility.

Cochlear modelling, and in particular whole (macroscopic) cochlear models, have played an essential role in the development of the travelling wave hypothesis. They have helped to refine, explain, and demonstrate the hypothesis since the travelling wave was first observed by Von Békésy. As a critique of alternative hypotheses, I think that if they neglect whole cochlea simulations, it will be to their detriment: not only does the scientific audience expect these type of tests, but they will also serve to refine and strengthen the hypotheses themselves. Chapter 7 shows an example of how a whole cochlea simulation can highlight features of the response that would otherwise be missed. A numerical cochlear modelling framework that could facilitate the implementation of these alternative models would lower the “barrier to entry”, particularly given that the alternative models are typically more complex to analyse in whole cochlear simulations, as this thesis will show. A framework of this type might encourage direct comparison between the travelling wave hypothesis and alternatives, which can only be good for the field.

1.2 Parallel Tracks

Two parallel tracks run through this thesis. In the first track we move from a discussion of the types of cochlear models we might require, to the specification of a general modelling framework¹ that fulfils these requirements, to a detailed description of what this specification entails, and ultimately examples of how it can be used. In the face of the diversity of uses for, and users of cochlear models, we draw on the anatomy and physiology of the cochlea to provide sensible limitations on the scope of

¹The source code for this framework has been released under the *Q Public License* (QPL), an open source license. It is available at <http://sourceforge.net/p/cochlear-emfw/>. All the other software required for this work is also available under open source licenses.

models that the framework should permit. Physics and mathematics inevitably also place restrictions on the types of models that we can consider within a single framework.

The material relevant to this track may be outlined as follows:

Chapter 2 A discussion of modelling in general and cochlear modelling in particular. The overall aims of the framework are set out.

Chapter 3 An introduction to, and overview of, the biological cochlea is given.

Chapter 4 Cochlear anatomy and physiology is used to inform the specification of the framework.

Chapter 5 The implications of the fluid-structure interaction that is present in the cochlea for the framework are studied.

Chapter 6 The details required to move from the general specification of the framework provided in chapter 4 to code that can be used to implement cochlear models are described. This code only provides part of the flexibility allowed in the specification.

Chapters 7 and 8 Two applications of the framework are given. The actual descriptions of the models' implementations follows a tutorial style. In each chapter background on the specific model considered is presented. In the case of chapter 7, this background includes further discussion of the arguments for and against the travelling wave hypothesis. (The work in chapter 5 also contains two further applications of the framework.)

The second track consists of a series of detours to investigate areas of difficulty for the proposed framework, and to establish key results that allow us to move forward. This track serves the development of our framework in two ways: it establishes a firm foundation for the framework in terms of results; and it validates central claims for the necessity of the framework, such as the claim that direct comparison between models in a single framework are preferable where possible. The results presented along this track have implications for the modelling community at large, and not just people who adopt this framework or the mindset behind it.

The material relevant to this track may be outlined as follows:

Chapter 5 Four important results are developed in this chapter. Firstly, the tools to analyse the *monolithic* model that underlies any cochlear model implemented with the framework are developed. Secondly, a characteristic feature of the fluid-structure interaction in the cochlea that can be utilised in the framework is noted. Thirdly, it is shown that the characteristic feature is also present for nonlinear models. This allows a general description of the implications of various types of nonlinearities for cochlear models to be given. Fourthly, an investigation into a possible partitioning scheme for use in the cochlear modelling framework begins to show

that the cochlear fluid-structure interaction is dominated by local interactions over short time scales. This is essential to the generality of the previous results.

Chapter 7 This chapter extends the final result from chapter 5. The nature of the fluid-structure interaction is further constrained, and the importance of interactions through the fluid over short spatial scales is demonstrated. In addition to completing the result mentioned above, this result is of general interest because the claim that the fluid does not favour these type of interactions is a common criticism of the travelling wave hypothesis.

1.3 Comments on Structure

A large number of cochlear models have already been proposed, and this thesis aims to facilitate comparison between existing and future models. The structure of the thesis has been tailored to meet its specific aims in various ways:

- Numerical results are primarily compared against published results for various models, rather than physiological results. The emphasis in this thesis is to demonstrate that the proposed framework is appropriate for a wide range of models. This makes replicating published results exactly the priority.
- The tone adopted in the thesis is generally of a “tutorial” nature. This is deliberate, because I hope to provide a document that is accessible to the full scope of potential users of cochlear models, in keeping with the aims of the thesis. I have kept the number of acronyms to a minimum in order to make the thesis easier to read for a general audience, preferring to write terms out in full. I have also provided background explanation on ideas such as state space that could have been briefer otherwise. However, I think that the style of writing that I have adopted will make the thesis more readable overall.
- Finally, I have used the first-person plural pronoun “we” in the majority of this thesis. Some people do this for stylistic reasons, to avoid using the first person pronoun. The reader is welcome to take my use that way if desired. However, I have used “we” as an invitation to the reader to join me in thinking about the topics that I am discussing. I have worked with other researchers on some of the subject matter in this thesis, so clarifying notes have been included where required to distinguish my own work.

1.4 Contribution

The aim of this thesis is to provide a modelling framework suitable for a variety of current and future cochlear models. For cochlear models that support the travelling wave hypothesis, the practice of

using numerical models to assess mechanisms is well established. However, a simple method of implementing new models and comparing them accurately to other models will certainly facilitate future work. Models that reject the travelling wave have *not* been assessed using numerical models to the same extent. The proposed framework can be used to remedy this situation. Rather than proposing a new model, this thesis proposes a new approach to cochlear modelling: an extensible framework in which (ideally) *any* cochlear model can be implemented. Ensuring this extensibility presents challenges on the same scale as proposing a new model, because the needs of the full spectrum of models must be considered when designing the framework. (By focusing on the physiology of the cochlea, and separating the concepts of *implementation* and *model*, overdesigning the framework is avoided.) Therefore, I think that the proposed framework, based on a broad survey of the present and future needs of the cochlear modelling community, and firmly rooted in the physiology of the cochlea, will be the lasting contribution that this thesis makes to the field.

However, I recognise that, with the long history of cochlear modelling and the breadth of approaches used, some might find the novelty of the framework difficult to assess. Three expressions derived to support the development of the framework provide a separate and clear contribution to the field. Equations (5.12–5.13) show the relationship between two approaches to cochlear modelling. Using these relationships, the pros and cons of three methods to implement cochlear models can be assessed. Furthermore, the relationship gives us greater freedom in deriving general results. Equations (5.9) and (5.23) show the formal relationship between linear and nonlinear cochlear models. By highlighting the terms in a linearisation that are affected by a nonlinearity, this expression allows a clear evaluation of the implications that various types of nonlinearities have for cochlear models. Finally, equation (7.20) formalises the nature of the fluid-structure interaction in the cochlea with a clarity and simplicity that has been lacking (hence, certain arguments concerning the fluid coupling in the cochlea persist). Therefore, both “track one” and “track two” mentioned above make an important contribution to the field of cochlear modelling. The contribution in track two is easier to assess, but the contribution in track one will be more enduring.

1.5 Related Publications

Over the course of this project, I have been first author or a co-author of a number of publications. In this section my contribution to these publications and their relationship to this thesis is summarised:

- [Rapson and Tapson \(2010\)](#) “Capturing the low frequency cochlear impedance in time domain models: The role of viscosity” – I was first author on this paper and it represents my own work for the purposes of this degree. This investigation influenced the structure of the proposed framework. The slides presented at this conference are provided in appendix E.

- [Rapson and Tapson \(2011a\)](#) “Emulating the theoretical bifurcation diagram of a cochlear model in computation” – I was first author on this paper and it represents my own work for the purposes of this degree. After the conference, this work grew into the next two journal papers.
- [Rapson et al. \(2012\)](#) “Unification and extension of monolithic state space and iterative cochlear models” – I was first author on this paper and it represents my own work for the purposes of this degree. The key results of this paper and its implications for the cochlear modelling framework are discussed in chapter 5. The paper is reproduced in appendix C.
- [Rapson and Tapson \(2011b\)](#) “Investigations into time stepping methods for cochlear models” – I was first author on this paper and it represents my own work for the purposes of this degree. The key results of this paper and its implications for the cochlear modelling framework are discussed in chapter 5. The paper is reproduced in appendix D.
- [Rapson et al. \(2011\)](#) “Exploring the travelling wave hypothesis with a finite element, state space, time domain cochlear model” – I was first author on this paper and it represents my own work for the purposes of this degree. The results presented at this conference are covered in chapter 7.
- [Hamilton et al. \(2010\)](#) “Understanding the mathematics of hearing using electronic circuits” – The results and text of this paper were prepared by Tara Hamilton. My only contribution was to assist in formalising the description of the cochlear fluid model. Therefore, this paper does not contribute towards this degree. However, my ongoing collaboration with Dr Hamilton around this topic was part of the motivation for my investigation in [Rapson and Tapson \(2011b\)](#), and chapter 8 arose from our collaboration. The results produced by Dr Hamilton, which appear in that chapter, can be clearly distinguished.

1.6 Recommended Reading

It inevitably occurs when writing a thesis that we first learn a concept or idea from one source but by the time we come to write up the work we have found a more specific reference to the idea that must be cited instead. This is particularly true of review articles, which by nature direct their reader to other sources. Both to acknowledgement of the input I have received, and as service to future researchers, I would like to list some references that I have found useful but had insufficient reason to cite otherwise:

- [Lyon and Mead \(1988b\)](#) and [Watts \(1993\)](#) were among the first works on cochlear modelling that I read. They provide a thorough introduction to the Liouville-Green² (LG) approach

²Max Viergever provides reasons for his preference for the name Liouville-Green in his PhD thesis ([Viergever 1980](#)).

(otherwise known as the WKB approximation). This method is a very important technique in cochlear modelling. We will refer to it a number of times in this thesis, but only in passing.

- [Ulfendahl \(1997\)](#) provided me with a good introduction to cochlear physiology and experimental methods. [Pickles \(1988\)](#) also served this role and has been cited elsewhere.
- In terms of cochlear modelling, I must also mention reviews by [de Boer \(1980, 1984, 1991\)](#) and [Allen \(2001, 2008\)](#). More recently, I had the opportunity to proofread [Duifhuis \(2012\)](#), and can strongly recommend it.
- [Fish and Belytschko \(2007\)](#) provides a readable introduction to the finite element method, and a thorough understanding of that material is more than adequate background for this thesis.

Chapter 2

Models

The problems presented by the study of tone perception lie entrenched in the debatable ground between several of the great divisions of natural science. They belong as much to physics as to physiology. In advancing to attack them, the ground must be tested at each step, to see if it be firm, and that ground belongs to the domain of anatomy. Should the attacker wish to push his advance by the means of mechanical aids, such as working models of the cochlea, he will find many small technical obstacles to be surmounted. He is not likely to advance far without the help and co-operation of the natives of the districts through which his operations are conducted, who know the country:—the physicist, the physiologist, the anatomist, and the practical engineer.

—George Wilkinson ([Wilkinson and Gray 1924](#), page v)

2.1 What is a Model?

2.1.1 General Usage of the Term

The word model is used in a wide variety of contexts in everyday speech and the sciences. It will be used in a very narrow sense in this thesis, even more narrowly than is common in the cochlear modelling field. Therefore, its current usage will be defined and justified. Before doing this, however, the wider uses in general and in the cochlear modelling field will be acknowledged by a series of examples.

In architecture, scale models that emulate the proportions of a finished building are used to give a sense of what it will be like. In this sense, a “cochlear model” might be a physical representation of the shape and form of the cochlea that could be used to teach its anatomy. While there would be certain similarities to the real cochlea, in particular relating to its appearance, at most points this type of cochlear model would have little in common with the real cochlea. It may seem strange to refer this type of physical model as a cochlear model, but the quotation at the start of this chapter shows




Figure 2.1: Mechanical cochlear model constructed by George Wilkinson, showing a view from in front and above (a) and a longitudinal section (b). (Adapted from [Wilkinson and Gray 1924](#), appendix I.)

that, in the early days of cochlear research, mechanical models dominated. Figure¹ 2.1 shows a 1920s cochlear model that consists of a brass box divided into two chambers with a flexible “basilar membrane” in the division. The models used by Von Békésy to demonstrate the travelling wave provide another example ([von Békésy 1960](#)). Interestingly enough, physical models are experiencing a renaissance as Micro-Electro-Mechanical Systems (MEMS) techniques allow one-to-one scale modelling of the cochlea. However, the emphasis in all these examples is shifted from providing a visual likeness to a behavioural likeness.

In modelling a behavioural likeness, analogy sometimes is taken even further. For example, the resonant behaviour of a mass-spring-damper system is similar to an inductor-capacitor-resistor circuit. This similarity is well known and has been used widely in cochlear modelling. [Wegel and Lane \(1924\)](#) proposed a cochlear model that consisted of inductors and capacitors. Figure 2.2 on the next page shows a network built by Bogert in about 1950 and a chip built by some of our collaborators in 2008.

Physical models of the cochlea were originally used to solve *mathematical* cochlear models that were too complicated for a closed form solution. Mathematical models express the behaviour of a system in terms of abstract equations. The majority of cochlear models have a mathematical model at their heart that is simply implemented in other forms. Many computational models, which now provide a viable option to solve these models, are also a way to study the behaviour of mathematical models. There is an alternative: a model may be constructed from basic units that inherently have similar properties to parts of the cochlea. If these parts are properly arranged their composition might show similar behaviour to the cochlea, despite the fact that precise mathematical descriptions have not been required. To some extent, the physical models discussed could be viewed in this way as

¹I am advised that reuse of published figures in a thesis falls under the “fair dealing” clause in the South African Copyright Act. Therefore, permission to use the figure (beyond the usual practice of citing the source) is not required. Where possible, explicit permission was nevertheless obtained.




Figure 2.2: Figure showing a 175-section network model of the cochlea from 1951 (a) and a 83-section model on a single silicon chip from 2008. (From [Bogert 1951](#), figure 1, and [Hamilton *et al.* 2008](#), figure 20.)

well. Computational models can also be built using this paradigm: adequate computational models of the various required parts can be combined in much the same way as their physical counterparts. Models that do not have an explicit mathematical description can serve two purposes: they can provide validation of mathematical models; and a means of signal processing that is similar to the cochlea.

The cochlear models discussed so far have all focused on describing or matching the mechanisms at work in the cochlea. Some models focus on obtaining the correct output for a given input, and interest in the underlying process is secondary. Statistical models are an example of this. The simplest relationship between the outputs and (minimum number of) inputs that provides the required accuracy is generally sought. The MPEG-1/2 Layer-3 (MP3) standard for audio compression makes use of a perceptual model based on psychoacoustics ([Brandenburg 1999](#)). This is an example of an auditory model that focuses on the output rather than the process involved. The perceptual model contains an implicit cochlear model because the cochlea is part of the auditory pathway it models, and many important parameters required by the perceptual model depend on the cochlea's behaviour.

In this discussion a number of distinctions have been drawn. Models that focus on the form or structure have been contrasted with models that focus on the behaviour. Sometimes a minimal model is required, and at other times great detail is desired. The idea of an abstract model and its implementation has been introduced. In the next section, this topic will be discussed further as we consider where to place the biological cochlea, and our models of it, in this paradigm.




Figure 2.3: Diagram showing an equilateral triangle (a), in which the mathematical idea of the object is emphasised, and two approximations to this shape (b) and (c). Both of these approximations use 13 regular objects in which the size of each object is allowed to vary. In (b) squares are used as the base object; in (c), circles.

2.1.2 Abstract Models and Implemented Models

In this section we begin by considering a simple example based on triangles. These prove to be a useful analogy to the process of cochlear modelling (and many other forms of modelling) in a number of ways.

2.1.2.1 The Triangles

Suppose that we require an object shaped like an equilateral triangle for a particular purpose. Equilateral triangles are first and foremost an abstract mathematical idea. They have a precisely defined length of side (assume it to be l) and angle between the sides (60°). Further properties can be proved, for example that the area is $l^2\sqrt{3}/4$. Assume that using these properties it can be shown that an equilateral triangle is the optimal shape for our purpose. However, no object in the real world can be a perfect equilateral triangle. For example, the edges will inevitably only be straight to within some tolerance. Fortunately, we are probably able to make an approximation to the ideal shape that will be adequate.

Suppose we wish to approximate the triangle by covering its surface area with either squares or circles. Figure 2.3 (a) shows an unfilled triangle that represents our abstract mathematical description of the ideal shape. (b) and (c) show two attempts to approximate the ideal with other shapes. (b) uses 13 squares of various sizes and (c) uses 13 circles. The objects are not allowed to overlap with each other. In each case, the largest object that fit in the available space was added at each stage. Both of these approaches can be made arbitrarily close to the desired figure by adding further squares or circles, but neither can perfectly fill the triangle with a finite number of objects.

2.1.2.2 What is the Ideal Solution?

Let the triangle approximation in (c) represent the biological cochlea. An animal has a cochlea for a reason – it provides an advantage to the animal, and therefore provides a solution a problem (in fact, a

number of problems). The biological cochlea represents a delicate trade-off between a wide range of competing factors. These include the solution's performance, and its energy and space requirements for example. Given this delicate balance, it might seem above dispute that the biological cochlea is also the ideal solution (a). However, if we focus on a subset of the cochlea's properties, say its performance, then we can construct an argument that it is not.

Focusing only on performance, we would probably find that a small change to the cochlea could improve it. This suggests that there is some other form that offers the best performance that can be obtained by this type of small change. The biological cochlea then approximates this in some way. In the same way that an equilateral triangle may be the optimal solution to some problems in an abstract sense, an ideal solution for the cochlea that is also an abstract idea may also be a useful construct. Implementations could at best approach such a solution. Trying to find carry this argument further by imagining a series of small changes that lead to the best performance fails, because we may arrive at a local performance maximum, and not a global performance maximum. This discussion quickly leads into the fields of mathematical optimisation ([Lange 2004](#)) and Pareto optimality ([Price et al. 2005](#)). It will not be pursued further. Instead, we turn to our example of the eye to round off the discussion.

The retina in the eye contains photoreceptors that are sensitive to light. The human eye is highly developed and performs remarkably well. Therefore, it is surprising that the photoreceptors face *away* from the light in the retina. The part sensitive to light lies behind the rest of the cell, with the axons located on the side nearest the light. These axons exit the retina at the papilla of the optic nerve, which produces a blind spot. The brain fills in this spot so we are not usually aware of it, but its presence can be demonstrated easily. The photoreceptors are orientated in the expected direction in the octopus. This suggests that these features in the retina are not necessary ([Dawkins 1986](#), chapter 4). The retina, then, approximates some better form in which these "faults" are not present. Similarly, the cochlea itself may not be the optimal solution.

This discussion has established that it is sensible to speak of an ideal solution that is different to the biological cochlea. The relationship between this ideal solution and the biological cochlea is the same as between triangles (a) and (c). Our current cochlear model is also an approximation to the ideal solution, but the technology used to implement our model is fundamentally different to the biological cochlea. Therefore, triangle (b) can represent our cochlear model. Depending on the purpose we have for our model, the ideal solution that we are working towards may not be the same as the one the biological cochlea approaches. However, if the ideal solution(s) and the biological cochlea are sufficiently far beyond our current models, then a step toward the biological cochlea is essentially a step towards the (possibly abstract) ideal solution to our particular problem as well.

2.1.2.3 Is Our Target Always the Ideal Solution?

In the preceding section, we established that the idea of an ideal solution that both our models and the biological cochlea attempts to approach. This need not be the end goal of all cochlear models. If we are interested in the biological cochlea as an end in itself (say to design a treatment), then the fact that it approaches some ideal solution is of secondary interest. If we were to make a replacement for the retina, we might want to include the papilla of the optic nerve because it is present in the biological retina. Despite believing that it is not necessary for the ideal solution, we include it for compatibility with the body. In terms of our triangles, the problem now is to find a model (b) that best approximates the biological cochlea (c), despite the difference in technology available.

2.1.2.4 Comparing Candidates for the Ideal Solution

Finally, we reinterpret the triangles as an analogy for the process of comparing cochlear models. Assume we have a candidate for the ideal solution in an abstract form, and now want to test its properties. Let the candidate be triangle (a). If we implemented (a) in two different technologies (such as a mechanical model and an electrical model) the two would differ from (a) in the same sense as (b) and (c) do. (b) and (c) also differ from each other, in spite of the fact that they are both approximations of (a). This situation occurs in cochlear modelling when various means of implementing the same mathematical models are used. As long as it is recognised that (b) and (c) are equivalent in some sense, the differences can give us extra information about the candidate solution, which is a good thing.

Another situation that commonly occurs is that (b) and (c) which are compared are *not* derived from the same abstract model. In this case, as we compare the differences between the abstract models using (b) and (c) we are, in fact, observing differences due to the implementation as well. This is not a problem if our interest is strictly in the output of the model. If our focus is on a description of the processes involved, care must be taken.

To avoid the dangers of comparing different abstract models using different technologies, we could choose one technology and implement both models using the same method. However, each technology has the potential to approximate some abstract models better than others. For example, if the candidate in (a) were a square rather than a triangle, method (b) would be favoured and (c) would have little benefit. Hence, it is important to consider whether the technology selected can be fair to both abstract models.

2.1.3 Summary

Models are used for a variety of purposes. A wide range of modelling approaches have been applied to the cochlea. This is partially because a single approach is not ideal for all uses, but also because

new approaches have become available. Conceptually, cochlear models aim to approximate one of two things to some level of detail and accuracy: the biological cochlea; or an abstract ideal cochlea². These purposes are distinct, but currently the biological cochlea is a sufficiently good approximation to any underlying ideal cochlea as to make the distinction moot. When implementations of abstract models are compared, any differences in behaviour observed are influenced by the method of implementation, as well as by the abstract models.

2.2 One Organ, Many Motives for Studying It

The preceding section has dealt with modelling as a general idea. It has talked about the types of models that exist and their categorisation for our purposes. In this section, the focus shifts to the users of, and the uses for, the models. The various reasons for studying the cochlea are discussed. Cochlear research is highly multidisciplinary, so the most relevant fields are summarised. Finally, the different requirements placed on cochlear models by the various uses are discussed.

2.2.1 Why Study the Cochlea?

2.2.1.1 Medical Knowledge and Treatments

A World Health Organisation fact sheet on “Deafness and hearing impairment” states the following estimate: “in 2005, about 278 million people had moderate to profound hearing impairment. 80% of them live in low- and middle-income countries” (WHO 2010). (Note that collecting this type of statistical data is difficult. Schein (2001) discusses some of the reasons that make it challenging.) The fact sheet also lists infectious diseases, excessive noise, injury, and the use of ototoxic (damaging to the cochlea) drugs as possible causes of hearing impairment. It states that half of the cases of deafness are avoidable.

This fact sheet shows that deafness affects a large number of people. The medical knowledge that we have accumulated to date allows a large proportion of deafness to be prevented. Seeing these cases prevented is as much a political problem as a medical one at this stage. However, further study of the ear can assist with increasing the number of preventable conditions, simplify the currently available preventions, and provide further treatments for people living with hearing impairment.

There are still open questions surrounding the structure and function of the cochlea. In particular, the outer-hair-cells are the subject of active research at the present moment. (Further details will be delayed until chapter 3.) Cochlear modelling assists in this research by providing a means to

²This topic is not just this author’s pet interest, it was raised in a question entitled “An ideal approach to cochlear modelling” at the International Symposium held at Titisee, Germany in 2002 (see Gummer *et al.* 2003, pg. 579–582). The discussion on that occasion covered some of the material presented in this chapter.

test theories developed in the course of the research. It also provides a link between studies of the structure (anatomy) and function (physiology). Models can be fitted to the anatomical data and their behaviour compared to the physiological data. Another area of active research is methods of delivering medication to the cochlea ([Salt and Ma 2001](#)). Models of the behaviour of the fluid in the cochlea can be used to study the diffusion of medication or ototoxic substances.

Two examples of modern treatments are the screening of hearing in newborn babies and cochlear implants. Babies can be tested for some forms of hearing loss within the first month of their life thanks to otoacoustic emissions (section 3.3.4.4, see also [Kemp 2008](#)). This allows early intervention that will assist in language development, if any problems are detected. Cochlear models are being used to further understand the otoacoustic emissions.

Cochlear implants are devices that include an electrode array that is implanted into the cochlea to stimulate the auditory nerve directly. The majority of the cochlear function is by-passed, so these devices can assist in situations that regular hearing-aids are unable to. Improved understanding of the cochlea and the technology used in the implants will improve the effectiveness of this treatment. Models of the cochlear geometry and the conductivity of the fluid have been used to study the array geometry and placement ([Hanekom 2001](#)). It seems reasonable to expect that an advanced cochlear model may perform initial processing for these devices, replacing the by-passed stages. Some testing of this idea has already been done with the technology available today ([Wolmarans 2005](#)).

Overall, there are exciting medical benefits to be reaped from continued study of the cochlea. Cochlear models of various types have played a role in understanding the ear and improving treatments. The evidence suggests that this will continue in the future.

2.2.1.2 Safety in Society and Industry

Our living environments are noisy places. There is concern about the level of exposure to high sound intensities in the workplace, at home, in the general environment, and during leisure activities. Strong evidence for a wide range of negative effects from noise exposure has been found. These effects include hearing impairment, hypertension, ischemic heart disease, sleep deprivation, annoyance, and reduced performance ([Passchier-Vermeer and Passchier 2000](#)). They were suspected in the 1960s and have been the subject of study since. In the opinion of Passchier-Vermeer and Passchier, the causal links between noise and these effects were adequately shown by the 1970s. They believe politics and policy are the main barriers to a reduction in the health effects of noise exposure. Therefore, although noise in our environment is a major problem, further cochlear research is unlikely to cause the policy changes that are required. However, details of the mechanism of hearing impairment are being researched and cochlear models can assist in this process. Various pharmaceuticals are being studied that may provide new avenues to reduce or prevent noise induced hearing loss ([Prell et al. 2007](#)).

2.2.1.3 Improved Presentation of Aural Information

Our understanding of the ear has shared a close history with the study of music. Hermann Helmholtz's 1862 work "On the Sensations of Tone" aimed to bridge the divide between acoustics and musical science (see [Helmholtz 1954](#) for an English translation). He achieved this with great success, and his work was still recognised as essential reading for both musicians and people working in physiological acoustics a century later. Musical ideas influenced physical and physiological research. This research in turn helped musicians to understand the science underlying their art.

The fields of research have since become more formal. Psychoacoustics, a branch of psychophysics, focuses on the perception of sound and its effects. That is, how a physical phenomenon (sound) becomes an internal perception. Details of the anatomy and physiology of the ear have been investigated, both separately from, and together with psychoacoustics (see chapter 3). Acoustics itself is now a wide field including both of these as disciplines, as well as a range of other topics such as ultrasonics and nonlinear acoustics.

Given the firm foundation that music theory currently has in acoustics, it is unlikely that new discoveries will have as much effect on the training and practice of future musicians as the ones to date have. However, there are a number of new applications of acoustics that stand to benefit greatly. The idea of acoustic virtual realities dates from the 1920s, but the idea has taken off since the 1990s ([Vorländer 2008](#)). This field is important for architecture, where it is used to simulate the acoustics of a building at the design stage. The entertainment industry also stands to benefit from improvements, in particular in computer games. In virtual reality scenarios, aural cues can provide information about the simulated environment such as its size and the location of objects. In the future this type of technology is likely to find much wider application. However, extended exposure leads people to become familiarised with these environments, and artifacts that were first unnoticed become apparent. This provides a demand for an ever more realistic experience ([Brandenburg 1999](#)). The higher degree of signal processing that can be done in these settings (compared to a musician's performance) allows algorithms to take advantage of a wider range of psychoacoustic phenomena. Hence, further research into the ear can result in improvements in the presentation of aural information.

2.2.1.4 Inspiration for Technological Advances

Studying the examples in nature can lead to inspiration and ideas that solve engineering problems. The fastener *Velcro*[®] was originally inspired by burrs of burdock that stuck to animal's fur. The inventor, George de Mestral, found some of these seeds stuck to his clothes after a trip in the Alps ([de Mestral 1952](#)). He looked at the seeds under a microscope to see how the mechanism worked, and invented a material along those lines. Birds have also long stood as inspiration to humankind to attempt to fly. The initial efforts attempted to copy birds down to details of flapping wings. Leonardo da Vinci considered designs of this type. They were not successful but, with the profile of aeroplane

wings emulating birds and combustion engines for power, aeroplanes are now able to fly many times faster than birds. However, in terms of efficiency and manoeuvrability we still have much to learn! In chemistry, spider's silk has inspired and continues to inspire advances in man-made fibres (Vollrath and Knight 2001). These well-known examples show the extent to which humankind inevitably adopts technology from nature.

There are certain technologies that it is natural to expect to benefit from a fuller understanding of the auditory pathway, and in particular the cochlea. These include: speech and speaker recognition, and audio compression and transmission, together with the downstream technologies such as communication systems, robotics, speech-to-text systems, etc.

Speech and speaker recognition technology generally consists of a *front end* that converts an audio stream into a series of basic symbols such as formants and pitch. A wide range of mathematical techniques are applied to extract these features. Some examples include Fourier transforms, the cepstrum, autocorrelations, and linear prediction. The main focus is to find a close match to the acoustics of the vocal tract, because this is where the signal originates. However, the idea that the ear performs a frequency analysis has been around since Helmholtz in the 1860s. Therefore, frequency domain techniques are inspired by the ear even though they draw heavily on the Fourier transform. There is evidence, and a general expectation, that human auditory system is well adjusted for range of auditory tasks, for example robust speech recognition (Kollmeier *et al.* 2008). The Ensemble Interval Histogram model is an example of a front end inspired by cochlear physiology that shows good robustness to noise (Rabiner and Juang 1993, section 3.5). Lyon *et al.* (2011) describes a range of biologically inspired “auditory filters” that can be applied to human and machine learning. However, physiologically inspired models are not yet widely used in these fields. Some of the reasons include complexity of the models, which results in significant latencies and high computational load (Kleijn 2008). In the human ear, the cochlea is an integral part of this front end. Therefore, further understanding of the cochlea may inform the otherwise separate advance of these technologies, as it has in the past. Equally, if an efficient physiologically accurate cochlear model is made (possibly in the form of a standalone silicon chip), this model may be incorporated into a front end in a production system.

Data compression is also a field that uses a wide variety of mathematical techniques. In the case of audio compression, we are not interested in the data itself, but rather in the perception that will arise from the data. This allows standards like MP3 (mentioned above) to perform lossy (the original audio cannot be recovered) compression that is more efficient yet causes the same perception. The perceptual model used in MP3 is based on psychoacoustics (defined below). A number of key features of the way we perceive sound are determined by the cochlea. (This will be discussed further in chapter 3.) Therefore, cochlear research may result in improved perceptual models for these types of systems, either directly through incorporating a cochlear model, or indirectly by suggesting modifications to existing models.

A final way that studying the cochlea inspires technology is through the demand it places on the research tools available. Robert Galambos observes that “the key to reporting biological information not already known is to use a new animal preparation, a new recording technique or a new measuring instrument” (Galambos 2006). Thus cochlear research, like any biological research, drives advances in the technologies that facilitate it. Cochlear modelling also challenges the technologies that enable it. These have historically been: first scale model construction; then computation facilities and digital signal processing; and more recently Very Large Scale Integration (VLSI) and MEMS technology.

2.2.1.5 Curiosity

Curiosity is a basic characteristic of humans, as well as other animals. When there is something that is new or unfamiliar we are inclined to investigate, even if it requires going out of our way. We often do not get any direct benefit from this activity. However, sometimes this curiosity pays off in a tangible way, either immediately, or at some point in the future. We are curious about the ear because it can do things we do not currently understand. As we investigate further we are likely to understand more, and at some level this is satisfying in and of itself. The investigation may surprise us by yielding results that benefit us in ways that cannot be logically expected at this point in time. Thus, by including curiosity as a valid reason to study the cochlea, I include the many unrelated benefits that can arise from studying something so intricate.

2.2.2 Differing Requirements Each Use Places on the Models

Considering the different uses for cochlear models discussed above, and the variety of researchers with an interest, it is clear that a single model of the cochlea is insufficient. Rather, each model must trade-off a number of factors in order to be optimal for a given need. The optimal trade-off for a given application will change as our understanding and the tools available improve. Some of the key factors in the trade-off are discussed below.

2.2.2.1 Detail versus Simplicity

This trade-off is intended to cover our choice of what to include in our model, and how the factors should be implemented. Including an extra factor adds detail. Finding an approximation that reduces the number of factors to consider makes the model simpler. If our aim is to approximate the ideal cochlea (from the triangles example) then our interest, like the physicist’s, is to find the simplest representation. On the other hand, if we want a close approximation of biological cochlea of a given species, details of that specific cochlea must be included. This is likely to include incorporating complexities that are not required for the ideal cochlea, but are necessary for its implementation using biological building blocks. Although in principle we may expect that a model of the ideal

cochlea should be simple and a model of the biological cochlea will be detailed, inevitably there is a trade-off between these ideas. As we approximate the ideal cochlea more closely, this too will require extra detail. Similarly, even models of the biological cochlea must compromise on the level of detail used because detail adds complexity.

2.2.2.2 Accuracy

Accuracy is a separate trade-off to the detail versus simplicity trade-off discussed above. That discussion focused on the choice of model, which has implications for accuracy, but accuracy can also be controlled separately to some extent. Accuracy is a measurable quantity, based on errors, whereas the detail versus simplicity trade-off is more conceptual. Accuracy is affected by a number of factors including: the uncertainty in the physical parameters used; the implementation; in computer models, the numerical precision of calculations (in physical models, similar limitations will exist); and the number of parts assigned to modelling the cochlear partition (see chapter 4). Some applications require high levels of accuracy; for example if a numerical result will be processed further in a speech recognition system. In other applications only the general trends in the response are of interest and high levels of accuracy are not required. Different methods of implementing a model will offer different levels of accuracy, and different means to control it.

2.2.2.3 Time Constraints

Processing an input signal or simulating the behaviour of a system takes time. Some applications require a result in “real-time” and processing delays of a few milliseconds can be noticeable. Other applications can be done offline. Even in these cases, the expense of occupying hardware and waiting for results causes time constraints. Therefore, while time constraints always exist, there are situations where they are the main design factor and other situations where they take their place behind other considerations.

2.2.2.4 Form Factor

The physical nature of a cochlear model is important for a number of applications. For example, if a cochlear model were to be included in a hearing-aid or cellular telephone, its size, power requirements, and manufacturing requirements would be important. An access control system might already include a desktop computer that a computational model could run on. For academic purposes, a supercomputer may be available for the simulation. Once again, it is clear that a single approach to cochlear modelling is not suitable for all the possible applications.

2.2.3 Fields Relevant to Cochlear Modelling

A surprising variety of researchers are interested in the cochlea. This is because the operation of the cochlea poses interesting questions in a range of fields. Unfortunately, this also means that researchers in each field need a functional understanding of the related fields in order to conduct their own research. These fields are discussed below.

2.2.3.1 The Biological Sciences

The cochlea has a complex anatomy and physiology. Many details of the anatomy have been understood since the 1800s, but the physiology was highly speculative until the early 1900s. The cochlea is very fragile and inaccessible, which makes it difficult to study. So, while there is already a large body of knowledge to incorporate into cochlear models, new results from improved experimental techniques must also be expected. The biochemistry of the cochlea also plays a role in its action – it is particularly important for understanding the active processes that are now known to exist in a healthy cochlea.

2.2.3.2 The Physical Sciences

Acoustics is first and foremost a discipline in physics. The ear uses devices like levers, funnels, and resonant structures to process the signal that it receives. Two of the central features in the cochlea are the transmission of the signal through a fluid and fluid-structure interaction. Furthermore, active feedback mechanisms that exist in the cochlea enhance its response. Physics provides the means to describe and model all of these behaviours. Mathematics enables physics in turn, and mathematical techniques, such as the Fourier transform and the finite element method, are used to solve some of these models.

2.2.3.3 Psychology

Psychoacoustics is the study of the perception of sound and its physiological effects. This branch of psychology focuses on the link between the physical world and our impression of it. Correlates for a number of psychoacoustic phenomena are known to be present in the cochlea before the signal is transmitted to the brain. Therefore, psychoacoustics gives valuable information about the cochlea itself, and not just the behaviour of the ear as a whole. Neuroscientists specialise in the function of the nervous system and the brain. Neuroscience is also relevant to cochlear modelling because the cochlea provides input to the nervous system and receives feedback control signals.

2.2.3.4 Engineering

Engineering is essentially the application of physics, so the comments above apply here as well. However, an engineering research project tends to be more applications focused than a physics project. The particular expertise expected in an engineer is the ability to combine theory with real world problem solving. Hence, engineers were called on to facilitate the early cochlear models discussed above. The field continues to interest engineers. In particular, the signal processing performed by the ear and the control mechanisms present in it are of interest to electrical engineers. Recent VLSI models are an application of analog electrical engineering techniques to the problem of cochlear modelling. Similarly MEMS models are an application of mechanical engineering techniques.

2.2.4 Summary

Cochlear models have the potential to impact a wide variety of fields. Each field places different requirements on the cochlear model. For example, some fields require numerical accuracy and real-time operation, whereas others emphasise detail and the ability to capture specific phenomena. The knowledge that can be incorporated into cochlear models spans a wide range of fields. This is both an advantage because models can facilitate collaboration, and a disadvantage because it is difficult to get an adequate overview of all the fields.

2.3 Statement of Problem

This chapter has looked at the role of cochlear modelling today and into the future. A discussion of what modelling is, laid a foundation for a survey of the potential applications for cochlear models. Some important trade-offs that need to be made when designing a cochlear model were noted. Discussing the fields relevant to cochlear research shows the diversity of knowledge covered.

This thesis aims to describe an approach to physiologically accurate cochlear modelling that will facilitate medical research. By facilitating this research, it is hoped that *all* applications for cochlear models will benefit. In fact, most of the fields discussed stand to benefit from cochlear models that are based on an improved understanding of the physiology. The approach should be highly flexible, because we do not yet know the final form models will adopt. A wide range of ideas will need to be tested, so it should allow various candidate models to be compared, placing only minimal restrictions on them. It must be as fair as possible to a wide range of types of models.

The modelling approach should also allow specialists to contribute their knowledge freely, without requiring a full understanding of a new field (cochlear modelling). Cochlear models already provide

a way to combine knowledge from the various relevant fields, but there are inevitably limitations to the degree that interactions between fields can be ignored. Whenever part of a complex system is studied in isolation, important modes of behaviour for the system as a whole are lost (Senge 1990). By allowing simulations of the whole cochlea, models should assist in avoiding this. However, the modelling approach itself should avoid introducing new dependencies where possible. In chapter 5 it is seen that the fluid-structure interaction in the cochlea must be handled particularly carefully in this regard.

Cochlear models have generally been introduced with a view to matching specific data. Therefore, the underlying approach used is seldom discussed in terms of being flexible enough to cover a wide range of models. Similarly, the primary user of the model is generally the researcher who developed it, so the ease of another person adapting it is also not often discussed. However, a few researchers have approached these questions from a similar angle. With the benefit of hindsight it is also clear that some models have filled this role. A brief review of some of these models that served as extensible “frameworks” is presented below. After that some details of the proposed approach are sketched, based on the discussion so far. Discussing “what won’t be covered” sets the scope of this work, and finally formal definitions are adopted for important terms that are used in an unusual way in this thesis.

2.3.1 A History of Cochlear Modelling Frameworks

Broadly speaking, cochlear models consist of two main parts: a model of the cochlear fluid and a model of certain structures, predominantly the cochlear partition (see chapter 3). The first computational cochlear models were solved in the frequency domain, with the cochlear partition represented by an impedance. This leaves only a spatial problem to be solved. A wide range of techniques were applied to the problem, including the Liouville-Green (LG) approach (otherwise known as the WKB approximation, Steele and Taber 1979a), Green’s functions (Allen 1977), and the finite-difference method (Neely 1981a). These methods each allow certain cochlear models to be implemented. The Liouville-Green approach, which relies on an approximation, is the most limited in terms of the models it can implement, but is valuable for the insight it gives into the solution. Green’s functions are more flexible, and have been used extensively by some researchers (for example Matthews 1980). Green’s functions are less flexible than the finite-difference method. The finite-difference method, first introduced to the field in 1978 by Stephen Neely, was extended to include (linear) active mechanisms in his Ph.D thesis (Neely 1978, 1981b). It continues to be widely used. Interestingly enough, transmission line models that were first used in the earliest cochlear models (for example Wegel and Lane 1924; Zwislocki 1948; Peterson and Bogert 1950) continue to be useful for a range of applications today (Puria and Allen 1991; Ramamoorthy *et al.* 2010; Epp *et al.* 2010), showing the flexibility of this particular modelling approach. These methods may be viewed as some of the frameworks that enabled cochlear modelling in the frequency domain to progress.

[Allen and Sondhi \(1979\)](#) introduced a new way to simulate the cochlear in the time domain. This method solves the (augmented) fluid problem and the structure alternately to compute next time step. [Diependaal *et al.* \(1987\)](#) performed a systematic study of the time stepping algorithm, and made improvements to the one suggested by Allen and Sondhi. In particular, he found that the Runge-Kutta 4th order algorithm was most stable and efficient for cochlear modelling. The majority of time-domain cochlear models to date use a time stepping algorithm following these lines. Hendrikus Duifhuis, one of the co-authors on that work, provided a motivation for a generalised time-domain cochlear model, together with an initial description in 2000 ([Duifhuis 2000](#)).

[Elliott *et al.* \(2007\)](#) introduced a new approach to the time stepping. Their approach forms a single, large state-space system on which the Runge-Kutta 4th order algorithm is applied directly, rather than through a series of steps. They aimed to provide a framework for cochlear modelling. Other researchers have adopted it (for example [Moleti *et al.* 2009](#)). Its extension to a wide range of current models has also been discussed ([Sisto *et al.* 2010](#)).

Thus, a wide range of approaches to spatial discretisation have been taken, but two major approaches to temporal discretisation (time stepping) have been used. In fact, even these two different methods can be shown to be mathematically identical (see chapter 5).

2.3.2 An Object-Orientated Computational Modelling Framework

Objects are a concept in computing that were first used in the programming language SIMULA from 1966 ([Dahl and Nygaard 1966](#)). While the idea is not new, it underwent significant development in the 1970s and 80s, and became popular during the same period. Object-orientated programming has some general advantages ([Bouzeghoub *et al.* 1997](#)):

- It encourages a close correspondence between objects in the program and parts of the real world systems modelled;
- Understanding of the details of each object is not required to use and extend the object;
- The behaviour of individual objects can be tested thoroughly – encouraging robustness;
- Re-use and extension of existing code is encouraged – reducing development time;
- Objects provide opportunities for exploiting parallelism; and
- The performance of C++ code (an object-orientated programming language) is similar to C.

These benefits align closely with the needs of the modelling approach that this thesis aims to introduce. A correspondence between physiological objects in the ear and computational objects in the simulation will make the modelling approach more transparent to users. Encapsulation of code into

objects means that users can safely extend and manipulate objects without knowing all the details of their implementation. Run-time and development time are both important. Re-using objects reduces the development time. Parallelism and the good performance of C++ allow run-time to be reduced as well.

Therefore, the substance of the modelling approach proposed will be an open-source C++ library that is an extension to the deal.II library (Differential Equations Analysis Library, [Bangerth *et al.* 2007](#); [Bangerth and Kanschat 2011](#)) . Primarily, it provides C++ classes to implement the finite-element method. However, the programming environment provided by deal.II (including intuitive linear algebra methods, a range of helper classes, extended error handling, and support of a wide range of platforms) also facilitates development. deal.II is under active development, and is up to date with the latest concepts in the finite element method and parallel computing. Use of C++ templates allows dimension independent software. Recently, support has been added for Intel Threaded Building Blocks and p4est – a library that allows thousands of computational cores to be used effectively ([Bangerth *et al.* 2010](#)). Therefore, deal.II is able to support even the most ambitious cochlear models. deal.II has been used in over 250 publications, illustrating its general usefulness.

deal.II's publication record also shows that an extensible library is able to facilitate a wide range of research. Given the scope of fields that could benefit from the proposed modelling approach, an extensible library is required, rather than a single application.

The flexibility to implement any model is ensured in part by adopting the finite-element method as the preferred method for spatial discretisation. Choosing to use an object-orientated approach does not restrict the set of feasible models in principle. Therefore this platform should allow free comparison of cochlear models. However, further detail of the cochlea is required to take this discussion further, so it will wait until chapter 4.

2.3.3 Power and Flexibility versus Safety and Efficiency

In computer programming, there are differing schools of thought on whether a programming language should try to prevent programmers from making mistakes with their code by preventing operations that are likely to be errors, or simply interpret what the programmers write assuming that it is all as they intend. The first approach places limitations on what a programmer can achieve with a given language: the programmer may have valid reasons for attempting to do something deemed a potential mistake and prevented. The second approach would be preferred were it not for the fact that humans do make mistakes, and they can be costly if they are not caught. (Paul Graham provides an interesting discussion of programming languages in chapter 10 of [Graham \(2004\)](#), also see the appendix to chapter 13, on the power of a programming language.)

Computer science formalises and studies tools for creating systems, so it is not surprising that the difficulty that language designers face is similar to the one we face when trying to specify a cochlear

modelling framework. We too have various opportunities to impose restrictions on the framework that go beyond those imposed by the mathematics or programming language, with the aim of preventing users of the modelling framework from making what we deem to be obvious blunders. Unfortunately, this path assumes that we are better able to judge what a blunder is than all the future users, and it runs the risk of excluding particularly novel ideas that are developed. Therefore, as a choice, we favour power over safety through imposing extra restrictions on the framework. Safety can be obtained by encouraging users to follow good programming practices, and to only make use of certain features as a last resort. Therefore, we try to impose as few restrictions as possible on the modelling framework. Furthermore, where restrictions must be imposed, we note why they need to be imposed, and how they might be relaxed should a future model require it.

There is also a closely related trade-off between flexibility and efficiency in the modelling framework. Offering multiple options can (and almost inevitably does) lead to overheads that could be avoided by restricting the possibilities available. This is especially true when we consider the range of other trade-offs that a flexible cochlear modelling framework must support. Each type of cochlear model offers a certain scope for optimisation that is lost when the code must be written to support a wide range of options. While some loss of performance in a general cochlear modelling framework is likely, it need not be overwhelming. deal.II itself aims to offer flexibility while avoiding compromising efficiency. This is achieved through the use of C++ features such as templates: a practice that is carried over into the framework where possible. Generality can also be a benefit to the framework, because if the efficiency of a piece of code is improved for one model, others may immediately gain the benefit as well. Optimisation of computer programs is best done once the code is working as desired, and it should be based on profiling of the program's performance. Trying to anticipate where the bottlenecks in a code are likely to occur is known to be difficult.

Regarding both the power versus safety trade-off and the flexibility versus efficiency trade-off, offering an open-source library allows the user significant advantages. With access to the code, they can improve the efficiency of the program in bottlenecks, either by removing options not required for the specific model, or by rewriting portions that are identified as problematic. They can also add support for additional features, making the library more powerful and flexible. Therefore, an open-source library means that users are not stuck with any mistakes or choices that we make.

2.3.4 What Won't be Covered

The key idea proposed in this thesis is that an object-orientated cochlear model can meet the needs of the field. All the material presented is organised around this theme. Work is presented to demonstrate that the key features (such as its flexibility) have been achieved. Unlike most cochlear modelling theses, no new model is proposed, rather a new modelling approach. Therefore, this work will touch on a range of models, some of which will be developed in more depth than others. However, none

will be developed to the same depth as in the typical cochlear modelling thesis that argues for a specific new model.

It is hoped that future researchers with an interest in specific models will adopt the approach and library developed in this work for their investigations. The framework should be suitable for both controversial and mainstream models. This thesis will *not* come to firm conclusions about the validity of either group of models on the basis of this work. To do so would be unfair to the methods, given that relatively little time has been spent assessing any specific model. A researcher committed to a controversial idea would certainly take pains to ensure that their idea is shown in the best light, which is not feasible under the circumstances. Furthermore, it would contradict the aim of this thesis, which is to establish a platform for comparison that is acknowledged to be fair.

2.3.5 Definitions Used in this Thesis

The term cochlear model has been used in a wide range of contexts, as section 2.1 showed. Without concise terms with clear definitions, language can easily become confusing or clumsy. Therefore, in this thesis the term “cochlear model” will be reserved for the abstract mathematical description of the cochlea, unless the context does not allow for confusion. The term implementation will be used for a specific realisation of an abstract model. It should be observed that a specific abstract model in a sense describes a family of models, generated by the selection of tuning parameters available in the model. Each member of this family of models may then be implemented in a variety of ways.

The term model has now been rigidly defined so it cannot be used to describe the generalised implementation methodology (approach) developed in this thesis. The term implementation is also not suitable because it already has a meaning assigned to it. However, the term framework appears to have the connotations sought. It has already been used occasionally in the text above, and some other authors have applied it in a way that is compatible with its meaning in this thesis (see [Elliott et al. 2007](#)). Therefore, a general framework is sought in which a wide range of (abstract) cochlear models may be assembled. This is done through a process of providing the governing equations in suitable form, together with the necessary tuning parameters. It results in a specific (cochlear model) implementation.

2.4 Summary

The process of, and need for, cochlear modelling has been discussed from conceptual and practical perspectives in this chapter. This discussion motivates the contribution that this thesis aims to make: to provide a modelling framework suitable for a variety of current and future cochlear models. This thesis will address the double-barrelled question of what features abstract cochlear models that are

useful to auditory researchers might contain, and in what form the widest selection of these features may be accurately implemented. Based on consideration of the uses of the framework, its outline has been sketched: an extensible library in an object-orientated language (C++). However, further details of the cochlea – the subject of our interest – are required to fill in the picture. The next chapter, which focuses on the anatomy and physiology of the ear, will provide the required details. It will also provide a sufficient background for readers who are unfamiliar with the cochlea.

Chapter 3

An Introduction to the Cochlea

As it was illustrated in section 2.2, a wide range of researchers are interested in studying the cochlea and for a variety of reasons. The field benefits from this diversity. This chapter aims to provide readers who have little or no background in cochlear modelling or the biology of the cochlea with the essential facts they need to understand this work. The material in this chapter is drawn from a variety of reviews and theses. In particular, [Pickles \(1988\)](#), [Dallos *et al.* \(1996\)](#), and [Kessel and Kardon \(1979\)](#) were used heavily. The following theses were consulted for a general overview of how they handled certain topics: [Neely \(1981a\)](#); [Matthews \(1980\)](#); [Watts \(1993\)](#); [Flax \(2008\)](#); and [Hamilton \(2008\)](#).

This overview will inevitably touch on areas in the literature where there is disagreement. In these cases, the predominant view on the matter will be presented, although the opposing view will generally be acknowledged. Areas of contention that are within the scope of this work will rather be addressed in the section where they apply. This approach allows a coherent and balanced picture of the cochlea and its function to be presented at this point.

3.1 Location in the Ear

The ear consists of three parts: the outer, middle, and inner ear. While this thesis focuses on the cochlea, which is the part of the inner ear that is specialised to detect sound, it is important to discuss its input and output pathways. When interpreting physiological data, it is important to know where and how the input was presented in order to make accurate comparisons between the model and the data. Furthermore, nonlinear cochlear models can require the input pathways to be included. Similarly, comparisons between cochlear models and auditory nerve data have proved useful, and feedback from the brain to the cochlea is known to be present. Figure 3.1 shows the anatomy of human ear and indicates the structures discussed below.




Figure 3.1: Diagram showing the location of the outer, middle, and inner ear, as well as cochlear nerve. (Adapted (slightly) from [Kessel and Kardon 1979](#), page 106.)

3.1.1 The Outer Ear

The outer ear consists of the pinna and the auditory meatus, which terminates at the tympanic membrane. The pinna contains a resonant cavity called the concha that leads into the auditory meatus. The volume of the human auditory meatus or ear canal is about 1 cm^3 . The outer ear as a whole creates a pressure gain at the tympanic membrane which is 15 to 20 dB at its peak for humans, and is relatively uniform between 2 and 7 kHz. It also has a characteristic effect on the gain of sound frequency components that depends on the location of the sound signal. This assists with sound localisation.

3.1.2 The Middle Ear

The middle ear consists of a cavity (the middle ear cavity) containing the malleus, incus, and stapes (shown in figure 3.1). These three small bones are collectively known as the auditory ossicles. The cavity is separated from the auditory meatus by the tympanic membrane, and is connected to the nasopharynx by the Eustachian tube that equalises the inner ear pressure to atmospheric pressure during swallowing. The stapes is connected to the oval window on the cochlea, as described further below. The tensor tympani muscle and the stapedius muscle can modify the behaviour of the ossicles.

The role of the middle ear is to transfer energy from the auditory meatus into the cochlea efficiently. It is also able to protect the ear against damage by loud sounds, particularly low frequency sounds. In its role as an energy transfer mechanism, the middle ear is often characterised by a transformer ratio. The middle ear often includes a second order transfer function model of the ossicles in addition to the transformer ratio in cochlear models (for example, [Neely and Kim 1986](#)). The earliest attempts to model the middle ear consisted of multiple stages, modelling each of the auditory ossicles ([Flanagan 1962](#)); these were generally simplified for inclusion in cochlear models. However, it has been shown that the middle ear has a considerable delay (possibly associated with wave propagation on the tympanic membrane but other sources contribute, [de La Rochefoucauld et al. 2010](#)), and that including this delay improves models significantly ([Puria and Allen 1998](#); [Parent and Allen 2007](#); [O'Connor and Puria 2008](#)). In fact, a closer match to physiological data can be obtained from a model that includes delay together with a simple model of the ossicles, than from an elaborate ossicles model that neglects delay ([Dong and Olson 2006](#)). In the gerbil middle ear, the pressure gain and delay were approximately 30 dB and $25 \mu\text{s}$ respectively. This held for a wide range of frequencies in the middle of the gerbil's auditory range ([Olson 1998](#)).

Cochlear models require an adequate middle ear model because the middle ear not only shapes the signal reaching the cochlea, but also provides an acoustic load on the cochlea, and a path for signals that arise in the cochlea to propagate back out. [Kemp \(1978\)](#) detected such signals, and they are understood to arise due to the nonlinear processes in the cochlea. However, each of the modelling




Figure 3.2: The main ascending neural pathways of the brainstem with minor pathways omitted. AVCN, anteroventral cochlear nucleus; PVCN, posteroventral cochlear nucleus; DCN, dorsal cochlear nucleus; LSO, lateral nucleus of the superior olive; MSO, medial nucleus of the superior olive; MTB, medial nucleus of the trapezoid body; NLL, nucleus of the lateral lemniscus; IC, inferior colliculus; MGB, medial geniculate body. (From Pickles 1988, figure 6.8.)

approaches mentioned above maybe suitable for a given cochlear model, depending on the purpose of the model.

3.1.3 The Auditory Pathway in the Brainstem

The cochlea is the site where mechanical to electrical transduction occurs in the ear. The final outputs of this transduction process are neural action potentials in the cochlear nerve (indicated in figure 3.1). Nerve fibres that carry signals from the cochlea towards the brain are called afferent fibres. Figure 3.2 shows the major ascending pathways in the brainstem. The auditory nerve fibres branch after entering the cochlear nucleus, with one branch innervating the anteroventral cochlear nucleus, and the other innervating both the posteroventral cochlear nucleus and the dorsal cochlear nucleus. The output from the anteroventral cochlear nucleus ends at the ipsilateral (referring to the the same side) lateral nucleus of the superior olive and medial nucleus of the superior olive, as well as the same structures on the contralateral (opposite) side. Note that these structures, together with other nuclei, form the superior olivary complex. While output from the anteroventral cochlear nucleus to the superior olivary complex forms the majority of the output from the cochlear nucleus, there is also output from the posteroventral cochlear nucleus to the ipsilateral lateral nucleus of the superior olive, as well as from both the dorsal cochlear nucleus and the posteroventral cochlear nucleus to higher levels of the the brainstem. These higher levels include: the nucleus of the lateral lemniscus; the




Figure 3.3: Centrifugal pathways from the superior olivary complex to the cochlea. UCOCB, uncrossed olivocochlear bundle; COCB, crossed olivocochlear bundle. (From Pickles 1988 figure 8.1.)

inferior colliculus; the medial geniculate body; and ultimately the auditory cortex. Only the cortex and medial geniculate body lie in the forebrain, the inferior colliculus lies in the midbrain, and the superior olivary complex, together with the cochlear nucleus, lies in the medulla oblongata.

Note that signals from both the ipsilateral and contralateral ear are used at all stages of the ascending auditory pathway. There is even input to the cochlear nucleus from the opposite cochlear nucleus that is not shown in the figure. Thus, there are numerous stages of processing before the auditory signal ultimately reaches the cortex. As the input to this system as a whole arises in the cochlea, physiologically accurate cochlear models could ultimately be (and to some extent already are) used to generate this input.

However, the auditory pathway in the brainstem is not only of interest as a potential output from the cochlea. There are also efferent nerve fibres that provide input to the cochlea from the brain. These arise primarily in the superior olivary complex. They form the olivocochlear bundle that exits the brainstem in the vestibular nerve and ultimately rejoins the cochlear nerve. As figure 3.3 shows, the olivocochlear bundle consists of both crossed and uncrossed bundles (innervating the contralateral and ipsilateral cochlea respectively). A few fibres from the bundle branch off to innervate the cochlear nucleus. The target of the remaining fibres in the cochlea will be discussed in section 3.2.

Aside from the fibres in the olivocochlear bundle that innervate the cochlear nucleus, there are additional pathways between the superior olivary complex and the cochlear nucleus. Centrifugal fibres from the nucleus of lateral lemniscus and the inferior colliculus also terminate at the cochlear nucleus. Thus, lower stages of auditory processing can receive input from higher stages that can modify the final sensory input. The centrifugal signals are influenced by both the previous acoustic input (previous because there is inevitably a delay associated with the signal transmission) as well as other factors such as attention. Finally, certain centrifugal fibres provide input to the tensor tympani ten-

don and the stapedius tendon in the middle ear that are able to affect the input to the cochlea at the stapes.

3.2 Anatomy of the Cochlea

3.2.1 The Scalae

The mammalian cochlea contains three fluid channels: the scala vestibuli; scala media; and scala tympani. They are approximately 35 mm long in humans. These channels run parallel to each other, and are separated by the Reissner's membrane and the cochlear partition respectively. (The cochlear partition is sometimes loosely referred to as the organ of Corti complex, or simply the basilar membrane in cochlear models.) The channels are coiled – with three and a half turns in humans – into a structure resembling a snail's shell that is embedded in the temporal bone. The channels narrow towards the *apex* of the cochlea, where the scala vestibuli joins with the scala tympani at the helicotrema. The opposite end of the cochlea is named the *base*, terms derived by association with a triangle. In humans, the cross-sectional area of the scala vestibuli and scala media combined is about 1 mm². These channels are generally lumped together in models of cochlear mechanics, because the Reissner's membrane is very thin. The average cross-sectional area of the scala tympani is also about 1 mm². The channels contain fluid. The perilymph in the scala tympani and scala vestibuli is high in sodium. The scala media contains endolymph, a fluid high in potassium. Cells in the stria vascularis are involved in maintaining the endolymph composition. (Figure 3.4 on the following page shows these structures.) The composition of these fluids also differs in terms of their protein and amino acid concentrations. However, their mechanical properties (for example their density and viscosity) are generally assumed to be similar.

The cochlea forms part of the inner ear, along with the vestibular system that assists with balance. These organs are in close proximity. In fact, both the scala vestibuli and the scala media are in direct fluid contact with various parts of the vestibular system.

The cochlea is surrounded by the bony otic capsule that has flexible membranes located at two points. The stapes, which serves as the primary point of acoustical input to the cochlea, is attached to the first membrane – the oval window. The oval window is located near the point where the scala vestibuli joins the vestibular system. The stapes is held in place by the annular ligament, which contributes a significant stiffness to the input impedance of the cochlea. The round window is the other flexible membrane in the otic capsule. It is located in the basal region of the scala tympani, providing a barrier between the cochlear fluid and the middle ear cavity. It also provides a pressure reference and allows the nearly incompressible cochlear fluid to move in response to a stapes input.




Figure 3.4: Cross-section through a turn of the cochlea showing the three scalae and the organ of Corti complex. (Originally in [Fawcett 1986](#), figure 35.11. Adapted from [Pickles 1988](#), figure 3.1.)

3.2.2 The Cochlear Partition

A cross-section through a turn in the cochlea (see figure 3.4) shows the three scala already described. Next to the scala tympani lies the basilar membrane that spans the gap between the spiral ligament and the osseous spiral lamina, and contains stiffening fibres. These fibres lie in a predominantly radial direction, parallel to the section shown. (This makes it natural to divide the basilar membrane into “radial segments” for modelling purposes.) The most prominent structure in figure 3.4 is the cochlear partition that is shown in detail in figure 3.5. The cochlear partition includes the organ of Corti that rests on the basilar membrane. The cells of Hensen (and the cells of Boettcher and Claudius, not visible in the figure) also rest on the basilar membrane, and are included in the cochlear partition.

The focus of modern modelling efforts has been on the hair cells located in the organ of Corti, whereas the basilar membrane was the focus of early efforts. Two distinct types of hair cells are present in the cochlea: a single line of inner hair cells that run along the length of the cochlea on the inner side of pillar cells; and between three and five rows of outer hair cells (increasing towards the apex) on the outer side of the pillar cells (see figure 3.6). In humans there are approximately 3000 inner and 12000 outer hair cells. Both types of hair cells have fine “hairs”, the stereocilia, which are of the order of $4\ \mu\text{m}$ long in the guinea-pig. Outer hair cells in Guinea Pigs are about $50\ \mu\text{m}$ long.




Figure 3.5: Cross-section through the organ of Corti. BV, blood vessel; IPC, inner phalangeal cell. (Originally in [Ryan and Dallos 1984](#), figure 22-4. As adapted for [Pickles 1988](#), figure 3.1.)

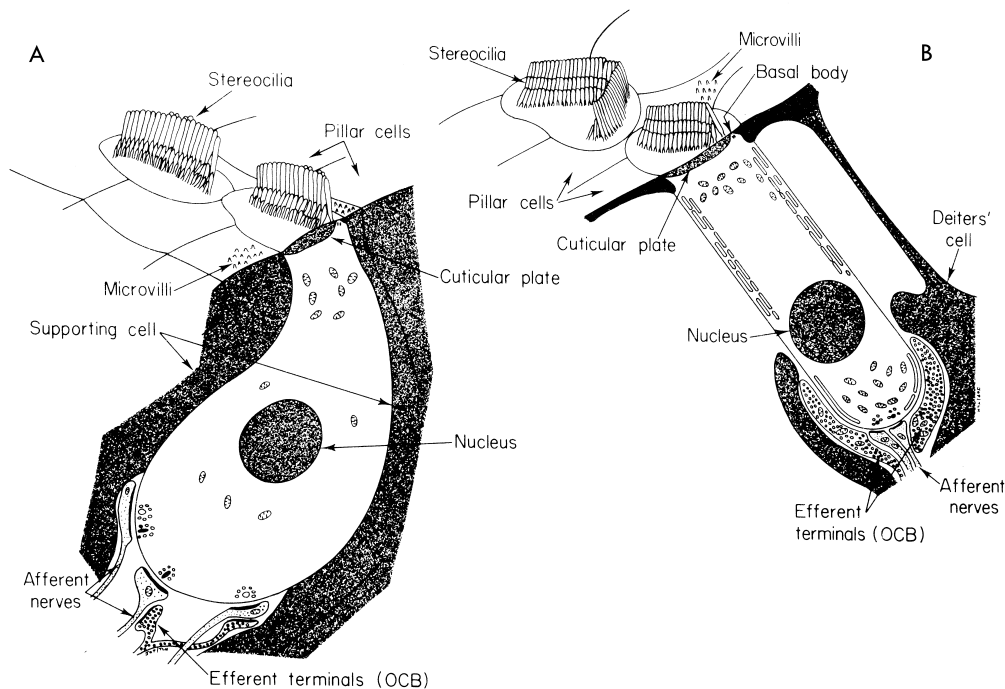


Figure 3.6: Diagrams of the inner and outer hair cells (A and B respectively) showing the characteristic shape of their cell bodies and stereocilia, as well as nerve synapses. (From [Pickles 1988](#), figure 3.5.)

The stereocilia of the inner hair cells form a row; those of the outer hair cell, a “v” or “w” shape. Deflection of the stereocilia affects gated ion channels that in turn affect the ion concentrations in the hair cells. The inner hair cells are surrounded by supporting cells, whereas the the outer hair cells are only supported by the Deiters cells at the end closest to the basilar membrane, referred to as the basal end. The opposite (apical¹) end of the outer hair cells form part of the reticular lamina, along with phalanges from the Deiters cells. The outer hair cells are able to modify their length and stiffness. Their method of attachment leaves them free to do so. The nerve supply to these cells will be discussed below.

It was originally thought that the organ of Corti was bathed in endolymph. Hence, the term “cochlear duct” refers to the scala media and cochlear partition, including the stria vascularis, Reissner’s membrane, and basilar membrane. In fact, the endolymph in the scala media is separated from perilymph in the scala tympani by the reticular lamina that forms the upper surface of the organ of Corti. This means that the stereocilia on the inner and outer hair cells are immersed in endolymph while the cell bodies are in perilymph. The organ of Corti is able to transmit movements of the basilar membrane to the reticular lamina and the stereocilia, yet it also allows some movement of the outer hair cells. The outer pillar cells are a key stiffness element in the organ of Corti.

The tectorial membrane lies over, but is separated from, the organ of Corti. It is described as gelatinous, although it also contains a dense network of fibres. The stereocilia of the outer hair cells are graded in length, and the longest stereocilia in each v-shaped bundle are embedded in the tectorial membrane. Whether the inner hair cells are also embedded in, lightly connected to, or unattached to the tectorial membrane is less easy to ascertain, and varies across species.

It has already been noted that the scalae taper, narrowing towards the apex of the cochlea. It is slightly counter intuitive, then, that the width of the basilar membrane increases from base to apex. In humans, the width increases from about 0.1 mm to about 0.5 mm. Similarly, the mass of the tectorial membrane increases, as does the size of cells in the organ of Corti. This variance causes a change in basilar membrane stiffness and cochlear partition mass with distance from the base. There are also size changes in the outer hair cells and length changes in their stereocilia with distance from the base.

3.2.3 Innervation

The cochlea receives its nerve supply and blood vessels through the modiolus, which is its conical central cavity. Figure 3.7 shows the main types of innervation that are present in the organ of Corti diagrammatically. The majority of the nerve fibres are afferent, about 30000 in man. Of these, about 95% innervate inner hair cells and are called type I cells. The majority of the remainder innervate

¹In hair cells, the terms basal and apical *do not* refer to the base and apex of the cochlea. In fact, the outer hair cells slant slightly towards the base of the cochlea from their supporting Deiters cell.

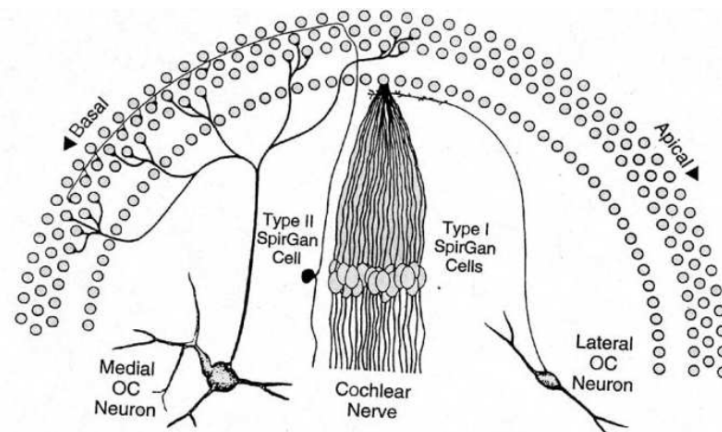


Figure 3.7: Diagram showing the innervation of the organ of Corti. (From [Warr \(1992\)](#), figure 7.12.)

outer hair cells and are called type II cells. These nerve fibres mainly use acetylcholine (ACh) as the neurotransmitter. A small percentage of the cells use a different neurotransmitter (γ -aminobutyric acid, GABA) and also innervate inner hair cells. There are multiple (approximately 20) afferent nerve fibres per inner hair cell, whereas single type II fibre forms synapses with between 6 and 100 outer hair cells. As discussed in section 3.1.3, the afferent fibres innervate the cochlear nucleus. The cell bodies of the afferent neurons lie in the spiral ganglion.

There are about 1800 efferent fibres that innervate the organ of Corti. These fibres arise in the superior olivary complex, near the lateral nucleus of the superior olive and medial nucleus of the superior olive. Neurons arising near the lateral nucleus of the superior olive have smaller cell bodies as well as small, unmyelinated axons. These axons make synapses with the dendrites of afferent nerves to the inner hair cells (figure 3.6A). Neurons arising near the medial nucleus of the superior olive are large, and have larger myelinated axons. These cross the tunnel of Corti (see 3.5 that indicates the pathway of axons in the organ of Corti) and form synapses directly with the outer hair cells. Figure 3.6B shows that the synapse surrounds the afferent synapse in this case.

Figure 3.8 on the next page shows the relative percentages of the centrifugal fibres arising both medially and laterally, and ipsi- and contralaterally. Notice that the majority of the efferent fibres to the inner hair cell are ipsilateral, whereas contralateral efferent fibres dominate the feedback to the outer hair cells. The afferent fibres all project to the ipsilateral cochlear nucleus. However, as noted previously, after the cochlear nucleus there are numerous nuclei that receive input from both sides of the brainstem.




Figure 3.8: Targets of the crossed (COCB) and uncrossed (UCOCB) olivocochlear bundle in the cochlea. Crosses indicate small cells near the lateral superior olivary nucleus that have small, unmyelinated axons and triangles indicate larger cells near the medial olivary nucleus that have large, myelinated axons. Some fibres omitted hence figures do not sum to 100%. (From Pickles 1988, figure 8.3.)

3.3 Physiology of the Cochlea

The preceding sections have discussed the location of the cochlea and its prominent structures. Although the anatomy of the cochlea does suggest possible functions for these structures, direct studies of their behaviour are required for confirmation. In fact, such studies have yielded numerous surprises over the years. In this section, the vulnerability of its behaviour will be discussed, followed by the passive behaviour that was studied first, and finally the active behaviour that was only discovered at a later stage.

3.3.1 Vulnerability

The behaviour of the cochlea's internal structures is difficult to observe directly. The cochlea itself is embedded in the temporal bone and surrounded by a bony wall. Experimental techniques that require access to the cochlea structures (for example to place a reflective bead for laser Doppler velocimetry measurements, or to measure the fluid pressure near the stapes) inevitably affect the state of the cochlea. The state of the cochlea can be assessed indirectly by recording the compound action potential in the auditory nerve, or cochlear microphonic potentials (due to the hair cells) at the round window.

Initially this was not well understood, so cadavers were used for Von Békésy's studies of the response of the basilar membrane to sound stimulation (von Békésy 1960). The frequency response of specific

locations were found to be broadly tuned. Using a light microscope, Von Békésy was able to observe the passive travelling wave in the cochlea for the first time. This behaviour was a surprise, because the travelling wave was not predicted by the dominant view at that time: Helmholtz's resonance theory of hearing. However, there were major discrepancies between the sharpness of the tuning in the observed response at the basilar membrane and the neural response of the auditory nerve. This initially led to the idea of a "second filter" that transformed the broadly tuned mechanical response to a sharply tuned neural response. In the 1970s, work by Rhode and his colleagues ([Rhode 1973](#); [Rhode and Robles 1974](#)) showed that the basilar membrane response was more sharply tuned than originally thought, but vulnerable to damage. In particular, the sharper response deteriorated rapidly after the death of the animal.

With due care, the cochlear response can be maintained postmortem, but experiments are generally performed on animals under anaesthetic. The animals are then "sacrificed" after the experiment. Even when working with anaesthetised animals, care must be taken in opening the cochlea, and some effect on the cochlear response must be expected and accounted for. Mats Ulfendahl discusses vulnerability further in [Ulfendahl \(1997\)](#), section 6.6.

3.3.2 Definition of Passive and Active Behaviour

The discussion of vulnerability, above, highlighted the fact that there is evidence for two modes of behaviour present in the cochlea. The "normal" behaviour of the cochlea is seen in living specimens that are in good physical condition, and a related, but qualitatively different, type of behaviour is seen in the cochlea after death. However, when tests that involve high input sounds are performed on an otherwise good specimen, the response of the cochlea can approach the response of a cochlea in poor condition.

In this section, a distinction between these two modes of behaviour that are generally referred to as the cochlea's active and passive behaviour will be drawn. Unfortunately, the cochlea's behaviour is not either strictly active or strictly passive, rather there is a continuum between these states, so care will be taken to discuss some of the implications.

The point at which behaviour should be defined as active is easy to select, but difficult to measure. The cochlea should be in perfect condition, in a healthy specimen, and the input should be at or near the threshold of hearing. Defining a similar point for the passive behaviour is more problematic. [Ruggero et al. \(1997\)](#) section III,F discusses this with reference to measuring the gain of the "cochlear amplifier". In principle, the passive behaviour of the cochlea is what we observe when the active (energy requiring) mechanisms present in the cochlea are not functioning, or are overwhelmed by the size of the input. It is the response of the passive structures present in the cochlea. However, in a dead animal the condition (for example, the stiffness) of even the passive structures changes. Figure 3.9 on the following page shows that the cochlear response 1 hour postmortem is

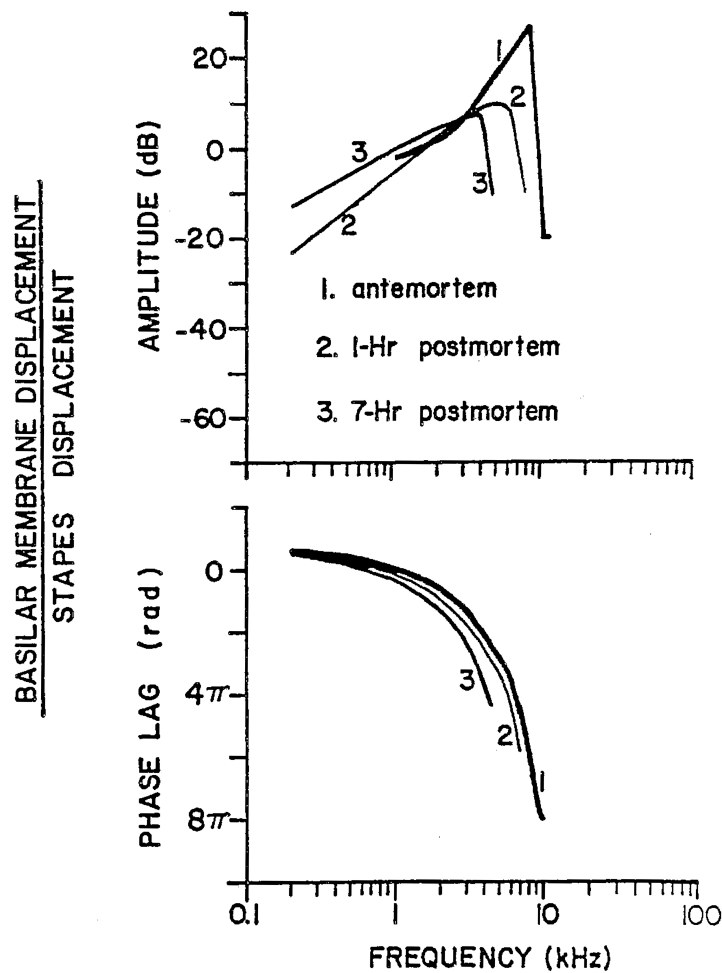


Figure 3.9: Diagram showing the effect of death on the basilar membrane displacement in a squirrel monkey. (Data from [Rhode 1973](#), animal 72-314. Figure from [Kim et al. 1980b](#)).

less sensitive and more broadly tuned than the antemortem response. Furthermore, the frequency of maximum response is shifted. After 7 hours the response has deteriorated further. Prolonged exposure to excessive sound pressure inputs can cause also permanent physical damage to the ear.

As noted above, researchers in the field are now aware of these problems, and are able to selectively (and even partially) remove the active mechanisms present: dead specimens can be maintained in good condition for the duration of an experiment; sound pressure levels can be selected to avoid damaging the cochlea; and drugs that selectively target the active processes can be administered. Nevertheless, these factors must still be considered, and corrections made, in particular with older publications, which are often still of value due to the difficulty and ethical concerns with collecting new data.

It is particularly important to be aware of the difficulties surrounding defining passive and active behaviour when dealing with cochlear models. Cochlear models span the entire period of modern


cochlear research. These models attempted to fit the best data at each point in time. Hence, for a period the models were attempting to fit what we now know is the nonlinear active cochlear response with models that are linear and passive. Von Békésy showed that the cochlear response postmortem is linear for a wide range of physiologically relevant stimulus intensities, so linearity of the passive response is generally assumed. A number of valuable cochlear models address matching the sharpness of the active cochlear response with a linear model. (This is due to the difference in computational requirements between a linear and nonlinear model.) Therefore, in the cochlear modelling there are two distinct antonyms used to describe a model: active or passive; and linear or nonlinear.

There is one further catch in the use of the terms active and passive referring to the cochlea. The term active properly refers to a system containing a source of energy. Hence, a system that is heavily damped, but has an active mechanism that reduces this damping by supplying energy to the system, would be called active. However, from a modelling point of view the system can be seen as a passive system that is not as heavily damped. It is common in cochlear modelling papers to reserve the word *active* for systems in which the active process puts sufficient energy into the system to cause a net increase in the local energy, even after considering the local loss mechanisms. [de Boer \(1996\)](#) suggests the term *locally active* to emphasise the fact that the net *local* energy is being considered. However, this term might be understood by some to imply that the activity is localised in some region of the organ of Corti, or even that the activity depends on the local behaviour only, as opposed to being directly influenced by more basal behaviour. In particular, the term *nonlocal* was coined in [Xin et al. \(2002\)](#), and has been used by some other authors when referring to feed-forward cochlear models (for example [Kim and Xin 2005](#) and [Sisto et al. 2010](#)). Therefore, there is a danger that readers who are not familiar with this history will take *locally active* and *nonlocal* as antonyms.


Having clarified these terms, the next sections will use the active behaviour versus passive behaviour paradigm to summarise some important physiological results.

3.3.3 Passive Behaviour

As noted in section 3.3.1, the first experimental evidence to shed light on the response of the basilar membrane to sounds was obtained by Von Békésy. The key results he found can be summarised in two diagrams, as shown in figure 3.10 on the next page. The top plot (a) shows the relative amplitude of the vibration of the cochlear partition with respect to frequency at six locations. Each location is most sensitive to one frequency – its so-called best frequency (the term characteristic frequency is also common, in particular when referring to the frequency that evokes the largest response in a given neuron). The figure also shows that the apical places are most sensitive to the lowest frequencies and basal places are most sensitive to higher frequencies. Part (b) shows the response of a section of the cochlear partition to four frequencies. The upper axis shows the envelope of the amplitude response, and the lower axis shows the relative phase response. In this case, it is seen that each frequency



(a)



(b)

Figure 3.10: Von Békésy's figures showing the frequency response of the cochlear partition for six location for constant peak stapes displacements (a), and the displacement envelopes and relative phase of the displacement for four tones of different frequencies (b). (From [von Békésy 1960](#), figures 11.49 and 11.58 respectively.)




Figure 3.11: Donald Greenwood’s original mapping of best place with respect to frequency. Solid points indicate Von Békésy’s experimental data. Labelled bars indicate regions that are supported by various other works. See the original (Greenwood 1961, figure 9) for further details.

causes a maximum amplitude at a particular location, called the best place. When a section of the cochlea was opened and viewed, Von Békésy described a travelling wave with a small amplitude at the base of the cochlea, increasing in amplitude up to the best place, and dying away rapidly after the best place.

The principle of a travelling wave in the cochlea is now the established view in the literature and among most researchers. It is also widely viewed as the basic behaviour that the active mechanisms enhance. This is disputed by certain researchers, however, who view the travelling wave as an epiphenomenon. This will be discussed further in chapter 7, where the transitory response of a passive model is studied, both to introduce the proposed modelling approach, and to discuss the response in light of this debate.

Von Békésy’s observation that the location of maximum amplitude depended on frequency was formalised by Greenwood, who fitted a function to the data available at the time (see figure 3.11). The result was significant because it provided a link between distance on the cochlear partition, frequency, and the psychophysical measure of critical bandwidth. This function is now referred to as the cochlear frequency map, and it provides essential data for fitting cochlear models. Passive models of the cochlea rely on a number of physical parameters, including the mass and stiffness of the cochlear partition. Using the cochlear frequency map, either the mass or the stiffness can be calculated from the other. This simplifies the process of fitting the model, and can highlight inconsistencies in the data. Data for the damping present in the cochlear partition are also needed. This is more difficult to measure than the mass and stiffness, so it is often estimated from the quality factor of the response to an input.

3.3.4 Active Behaviour

The active behaviour of the cochlea has been studied since the 1970s, when improvements in research devices allowed their effects to be observed. However, evidence that there are nonlinear processes at some point in the auditory system was already available from psychoacoustics – the study of sound perception. Therefore, research in this period established that the nonlinearities are present at the level of the basilar membrane motion (as opposed to neural nonlinearities), and that an active response in the cochlea is required. The evidence for nonlinear active processes in the cochlea came from a number of places, beginning at roughly the same time. While the researchers who first observed these effects initially struggled to get their results accepted, the increasing weight of the evidence eventually won the day.

3.3.4.1 Nonlinearity

In 1971 William Rhode used the Mössbauer technique to measure the vibration of the Squirrel Monkey basilar membrane (Rhode 1971). The Mössbauer technique uses a small radioactive gamma ray source, placed on the basilar membrane, to detect very small velocities. Rhode was able to detect velocities in the range 0.2 to $2 \text{ mm}\cdot\text{s}^{-1}$. Although the technique had been used previously to measure basilar vibration (Johnstone and Boyle 1967), Rhode was the first to observe nonlinear behaviour. Von Békésy had only described linear motion, and so Rhode was cautious in his claims – stating only that the system he observed (with the source in place) was nonlinear. His results were confirmed by other research groups and with another technique: laser velocimetry. Laser velocimetry is inherently linear whereas the Mössbauer technique is in itself nonlinear. This provided final confirmation that the nonlinearity observed was not an artifact due to the measuring technique.

Figure 3.12 shows some more recent results for the nonlinear behaviour of the cochlea taken from Ruggero *et al.* (1997). The data is taken from a single animal and is represented in three ways. Parts (a) and (b) show velocity-intensity functions, separated into frequencies above and below the characteristic frequency for readability. The response for frequencies near the characteristic frequency (10 kHz) grow at a slower rate than 1 dB/dB , which is a linear relationship, hence the response exhibits a compressive nonlinearity. Ruggero *et al.* note that the response near the characteristic frequency is compressive even at the highest stimulus intensities that are physiologically reasonable. Part (c) shows the velocity of the basilar membrane as a function of frequency for a specific stimulus intensity. Louder sounds cause a larger vibration, but the frequency selectivity of the response is reduced, and there is a shift in the best frequency too. Finally, part (d) shows the data expressed as gain with respect to frequency for fixed stimulus intensity. The shape of each curve is similar to (c), but this plot makes it clear that the cochlear response shows the highest gain at the lowest signal intensities. The average stapes motion is indicated in the figure. It is clear that stapes motion does not account for the nonlinear gain. In this figure, linear growth is indicated by lines that overlap. Hence, the

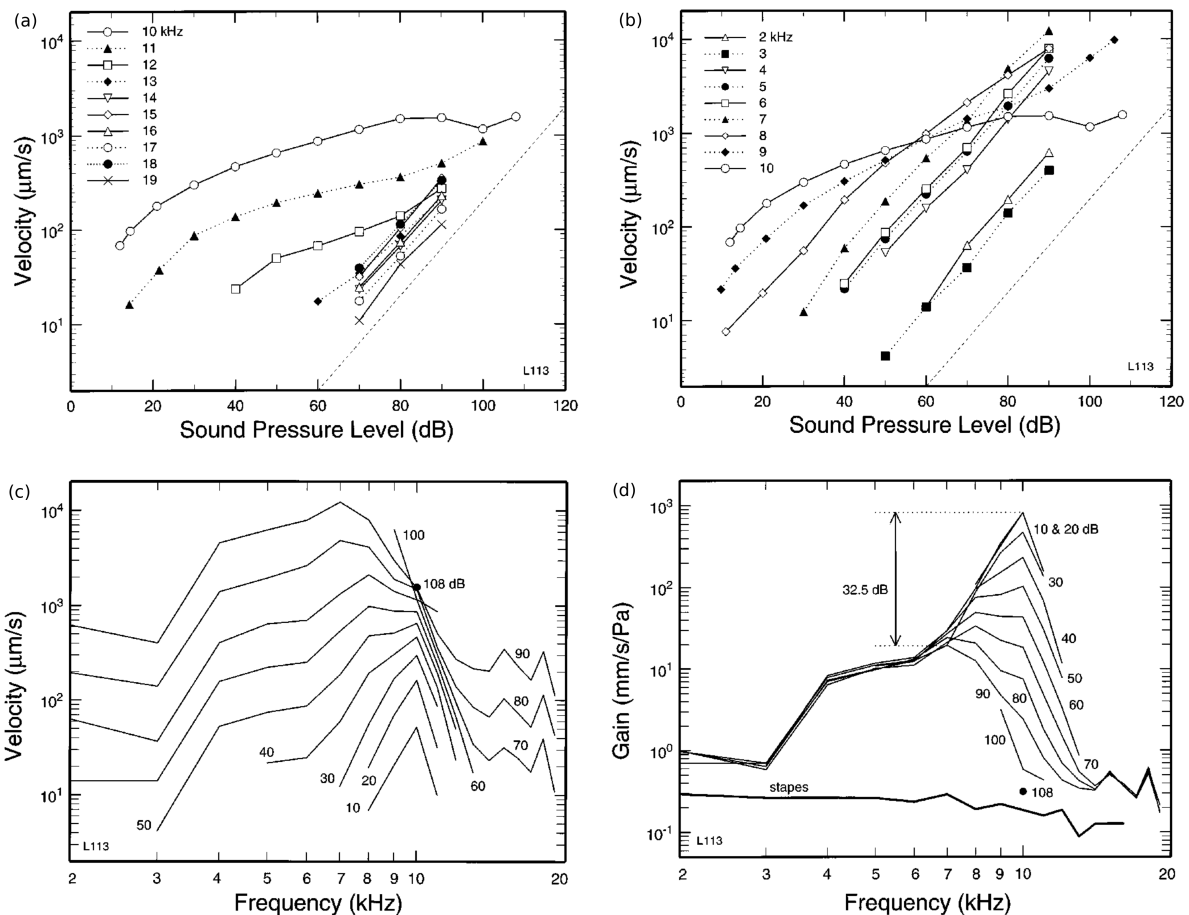


Figure 3.12: A collection of figures from [Ruggero et al. \(1997\)](#) that shows the type of data available to describe the nonlinearities in the cochlea. All the data is taken from a single animal. Three different means of representing the data are shown. (a) and (b) show velocity-intensity functions for tones below and above the characteristic frequency (10 kHz). The straight dashed line indicates a linear growth of 1 dB/dB. (c) presents the isointensity curves for the velocity of the basilar membrane response to tone pips of various frequencies. (d) shows isointensity curves for the gain (velocity divided by stimulus pressure) as a function of frequency. The average stapes motion is indicated by the marked line at the bottom of the figure. One definition for the gain of active response with respect to the passive response in this cochlea (32.5 dB) is indicated on the figure.




Figure 3.13: Data in figure 3.12 above plotted as isovelocity contours (solid lines marked 25, 50, 100, 200 and $400 \mu\text{ms}^{-1}$) and isodisplacement (dashed lines marked 0.8 and 0.4 nm). This representation allows comparison with the average frequency-threshold tuning curve. (From [Ruggero et al. 1997](#), figure 11.)

response is linear for stimulus intensities below 20 dB for this cochlea. [Ruggero et al.](#) describe a region of linear growth for small stimulus intensities. This result is in conflict with other views that the cochlea has “essential” nonlinearities – that the behaviour remains nonlinear even at the lowest audible sounds ([Eguíluz et al. 2000](#), but see their comments on page 5234 of that report).

Expressing the data as a tuning curve (figure 3.13) allowed [Ruggero et al.](#) to draw comparisons with neural tuning curves. They note that the tuning curves show a good match near the characteristic frequency, but the match is worse for much lower frequencies. This figure is included as evidence that there is a large degree of correlation between the neural and basilar membrane response, so it is no longer necessary to postulate a second filter. The idea of a “second filter” is still sometimes referred to in the sense of trying to find the process that causes this sharpening of the basilar membrane response from its passive to active behaviour. Essentially this amounts to couching recent ideas in the language of previous researchers for stylistic reasons.

3.3.4.2 Two-Tone Suppression

Two-tone suppression was first identified in the auditory nerve response. [Rhode](#) demonstrated it in the basilar membrane response as well ([Rhode 1977](#)). However, he used higher sound levels than required to observe it in the auditory nerve. Figure 3.14 shows more recent results for the effect.

The experiment studies the interaction of two tones, which would show superposition for a linear system. The strength of a probe tone, the frequency of which is near to the best frequency of the

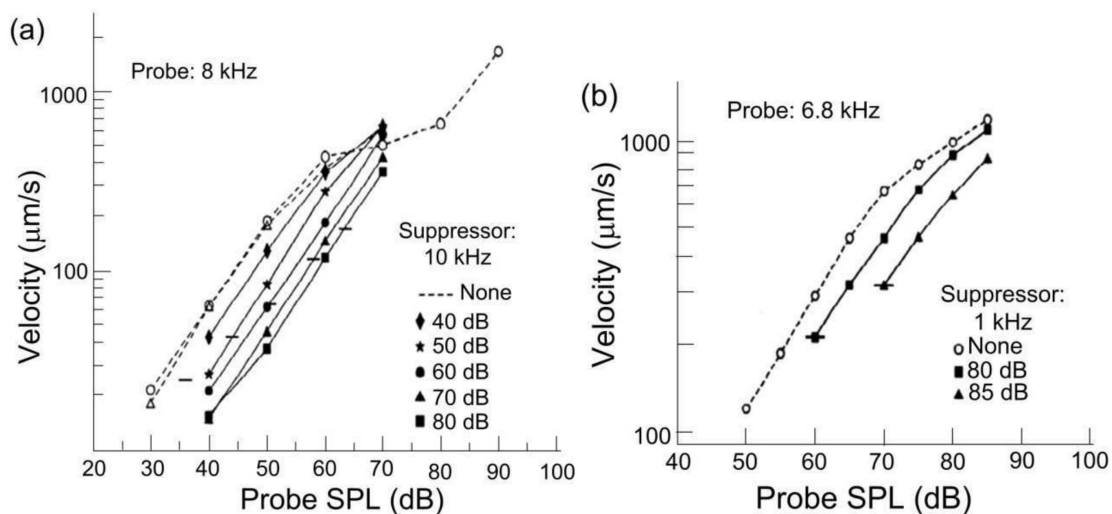


Figure 3.14: Diagrams showing the basilar membrane response to probe-tones near the characteristic frequency of the measurement location alone and with suppressor tones. (a) shows the response to a higher frequency suppressor, measured with laser velocimetry. (b) shows the response to a low frequency suppressor, measured with the Mössbauer technique. (From [Ruggero et al. 1992](#), figures 3 and 4. As appeared in [Hamilton 2008](#), figure 2.8.)

location on basilar membrane under study, is varied in the presence of a second suppressing tone. The suppressing tone modifies the effect of the probe tone, thereby showing nonlinear behaviour. The figure shows the effect of a higher frequency suppressor in (a), and a lower frequency suppressor in (b). In both cases the response of the probe alone is also indicated. In either case, the response to the probe and suppressor tones together can be smaller than the response to the probe alone. The magnitude of the suppression depends on the level of the probe tone, and increasing the suppressor tone magnitude causes an increase in suppression. Suppressors below the characteristic frequency show a higher rate of growth in suppression with suppressor intensity than suppressors above the characteristic frequency. Damage to the cochlea reduced the suppression, and the magnitude of the suppression obtained was correlated with the sensitivity of the basilar membrane response.

3.3.4.3 Distortion Products

Two-tone distortion products or combination tones are additional sounds that may be heard when two pure tones (f_1 and f_2 , with $f_2 > f_1$) are presented to a listener. The sounds match specific combinations of the original pure tones. $f_2 - f_1$ and $2f_2 - f_1$ are generally the most prominent, but combinations of the form $(n + 1)f_1 - nf_2$ and $(n + 1)f_2 - nf_1$, for n integer, are present at differing magnitudes.

These combination tones were observed directly on the basilar membrane only in 1991 by Robles *et al.* using laser velocimetry ([Robles et al. 1991](#)). Previous attempts to identify the distortion products were unsuccessful due to their small size, and the fact that the Mössbauer technique was used, which

is itself nonlinear. However, researchers had indirect evidence for the presence of distortion products in the basilar membrane response, such as the results presented below.

[Kim *et al.* \(1980a\)](#) inferred the distortion products on the basilar membrane from cochlear nerve data. Figure 3.15 shows a typical stimulus used for the experiments, including the frequencies f_1 and f_2 , the resulting waveform, and its Fourier spectrum. The frequencies were selected as harmonics of a fundamental, which ensures that all distortion products are also harmonics. In particular, $f_2 - f_1$ is the second harmonic, and $2f_1 - f_2$ is the fifth harmonic. First, the characteristic frequency of up to 418 fibres in the cochlear nerve were found, providing a mapping between fibres and their approximate place on the basilar membrane using their characteristic frequency. They then collected period histograms of the nerve fibres' response to the waveform, and examined the Fourier spectrum of this histogram for the distortion products.

The characteristic frequencies of the neurons presented in the figure correspond to $f_2 - f_1$, approximately $2f_1 - f_2$, f_1 , f_2 , and a frequency greater than f_2 . The two input frequencies and the two prominent distortion product frequencies are evident in the response. Furthermore, the response is strongest for each frequency in the nerve fibre tuned to that frequency, although other frequencies are present.

Evidence from experiments where the best place for f_1 and f_2 are temporarily, or permanently, destroyed by moderate, or intense sound respectively, suggests that the tones are generated in the region where both f_1 and f_2 have a strong response. When these regions are damaged the distortion products are affected, although the response to a tone of the same frequency as the distortion product is unchanged.

3.3.4.4 Otoacoustic Emissions

In 1948 Thomas Gold suggested that the cochlea must be active to account for the (overly) high quality factor that he thought was present in the cochlea ([Gold 1948](#)). While his views ultimately proved right in many ways, they were controversial when suggested. One of the reasons that they were controversial is that an active cochlea would seem to imply that there should be sound radiated from the cochlea. Gold argued that this would not be the case if the cochlea had a smooth transition in resonance frequencies. He expected to only see radiated sound when there was excess feedback at some point. Although he looked for this feedback, he was unable to demonstrate it using the apparatus available at the time and the approach he adopted.

However, indirect evidence for tones generated in the cochlea was available. There were unusual medical cases linked to tinnitus (often associated with hearing loss due to damage to active mechanisms) in which emissions could be heard from the ear. Psychoacoustic evidence such as beats (periodic increases and decreases in the perceived loudness of a tone when two tones of similar frequency are presented together) and combination tones that could be obtained with only a single input

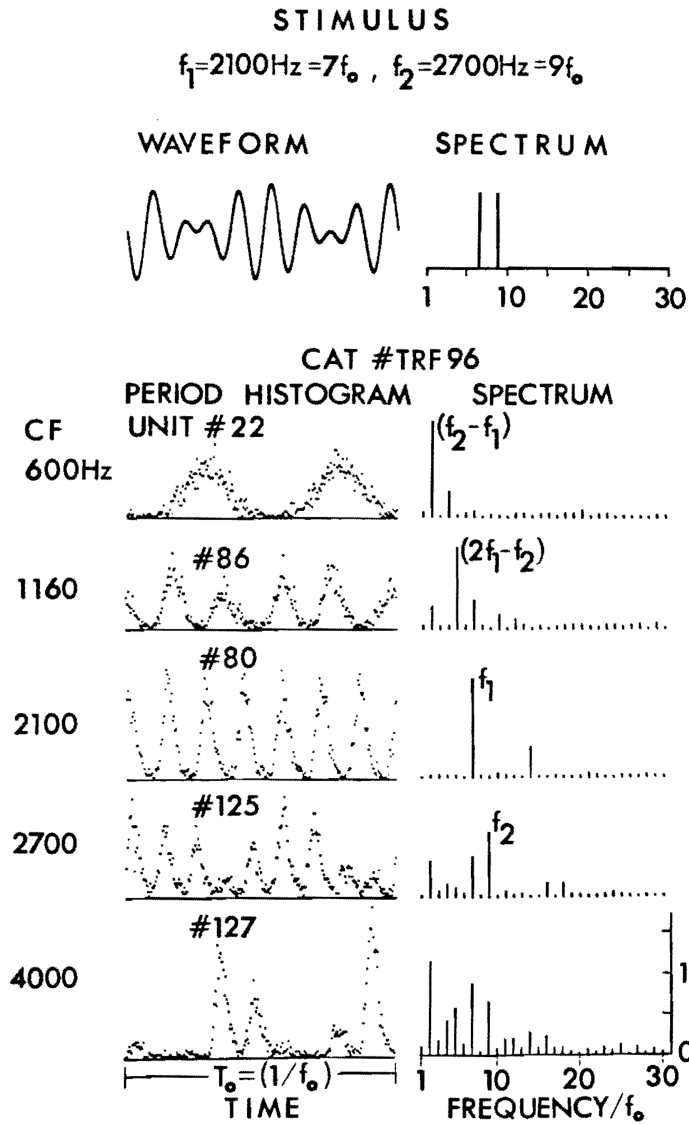


Figure 3.15: Figure showing distortion products on the basilar membrane, as inferred from period histograms from the cochlear nerve. The characteristic frequency of numerous nerve fibres was used to identify their location on the basilar membrane. Their response to two simultaneous, harmonically related, tone inputs was recorded. The stimulus frequency, waveform, and spectrum are shown. The resulting period histogram and spectrum are shown for various fibres. The spectrum shows prominent peaks at frequencies not in the stimulus. (From [Kim et al. 1980a](#), figure 1.)

tone suggested that a second tone was present in the cochlea.

David Kemp recorded the first otoacoustic emissions in the course of experiments to test a hypothesis about the source of the fine structure present in audiogram thresholds of healthy ears between 600 Hz and 3 kHz. Figure 3.16 shows some of the results that he presented in his 1978 paper (Kemp 1978). The top left panel shows an example of the time response to a click in a chamber of the size of an ear canal, compared to a click in a human ear. A delayed response is visible in the recordings of the ear. The top right panel shows evidence for cube-root growth of the response in the lower part of the image. The upper part of that image shows a frequency decrease with time that is also known to occur in the cochlea itself. The lower left panel shows further evidence for compressive growth of the response with increasing stimulus intensity. Finally, the lower right panel compares the response to clicks and tones. Similar size clicks cause different responses in different ears. (Results from four ears are shown in the figure.) The frequency of the tones is visible in the response. The results for the tones were obtained in a single ear.

Otoacoustic emissions are recorded in the ear canal by a sensitive microphone in a probe with a soft plastic tip to seal the ear canal behind the microphone. Kemp (2008) provides calculations to show that typical otoacoustic emissions require the ear drum to move just 0.01 nm, or one-tenth of the diameter of a hydrogen atom. He states that current instruments can measure vibrations of the ear drum of 0.3 pm. This means that otoacoustic emissions currently provide the most sensitive method of studying cochlear activity. Furthermore, these measurements are noninvasive, and can be used in humans. In fact, the original work was based on human ears, with David Kemp using himself as one of the subjects. Therefore, otoacoustic emissions provide a valuable source of information about the physiology of the cochlea.

A number of different ways of measuring otoacoustic emissions have been developed. These are: spontaneous otoacoustic emissions; stimulus frequency otoacoustic emissions; transient or click-evoked otoacoustic emissions; and distortion product otoacoustic emissions. Kemp (2008) provides an excellent summary of these different observation methods that is included as an abstract in this thesis because it is so succinct (see appendix A). However, the author makes the point that these should not be viewed as separate phenomena. Rather, these are the response of a single system to different inputs, and they ultimately provide information about that system.

3.3.5 Hair Cell Physiology

The previous section has discussed the behaviour of the cochlea as a whole. This includes its sharp tuning and nonlinear response. It also provided some details about a physiologically vulnerable active process that provides this behaviour, because we noted that death of the animal destroys the response. We now know that the active behaviour of the cochlea depends on the outer hair cells that

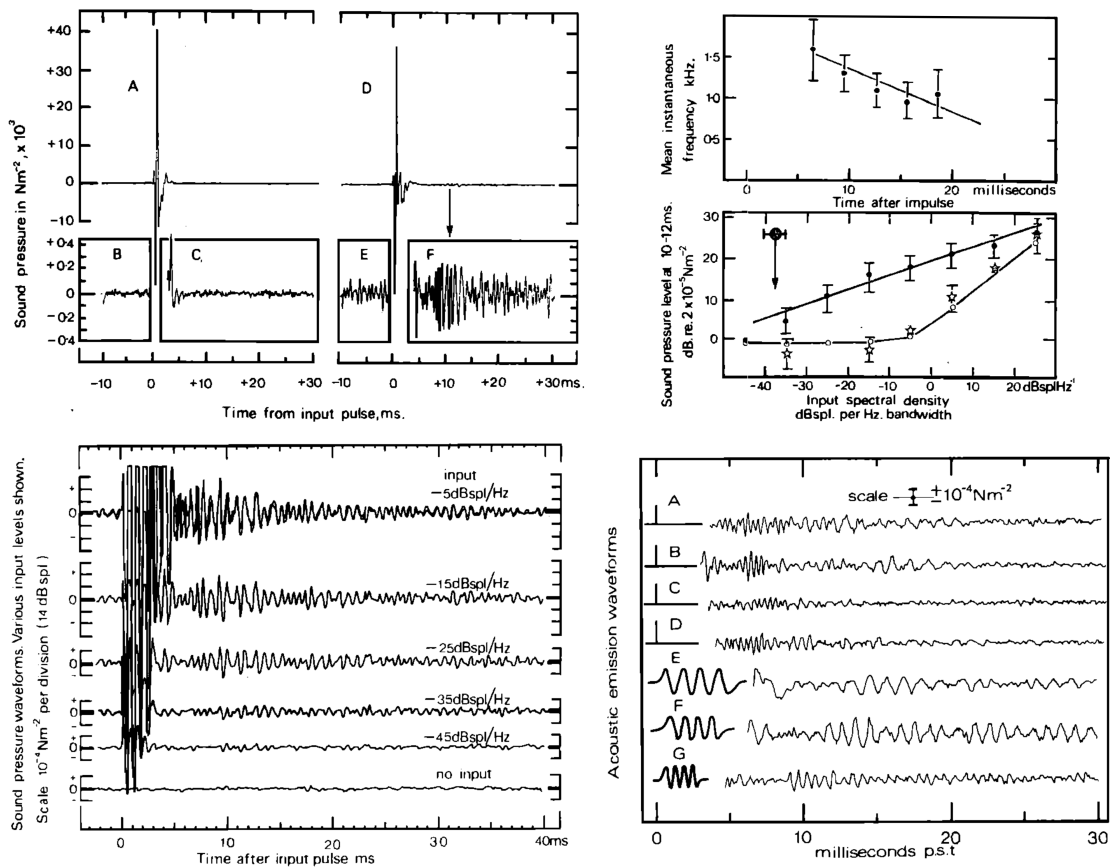


Figure 3.16: First evidence of otoacoustic emissions published in [Kemp \(1978\)](#). This image follows the layout in [Kemp \(2008\)](#). Image in the *upper left* compares a click in a sealed chamber of the size of the ear canal (A–C) to a click in a healthy ear (D–F). The sealed chamber shows only noise before and after the click (B and C show the signal magnified) whereas the human ear shows a delayed response that is a maximum about 10 ms after the original click and lasts more than 25 ms. The signal visible in E is response to previous clicks. This is an example of a *transient evoked otoacoustic emission*. The *upper right* figure, upper panel, shows the change in the frequency of the click-evoked emission with time. This provides evidence of the cochlear origin of the response because the cochlea is known to show a similar property in response to clicks. The lower panel of the same figure shows the growth of response with stimulus intensity. Filled points indicated data from healthy ears, unfilled circles the cavity measurements, and stars mark data from ears with cochlear deafness. The cavity measurements show a close match to the data from ears with cochlear deafness. The linear increase at the higher stimulus intensity is due to the input affecting the measurement. The response of the healthy ears shows a cube-root growth: as the stimulus increases by 60 dB, the response increases by 20 dB. *Lower left* shows the effect of stimulus intensity on transient evoked otoacoustic emissions by plotting the time response. Once again the response grows in a compressive manner as stimulus intensity is increased. The *lower right* shows some responses to click evoked otoacoustic emissions from different ears (A–D) and tone evoked otoacoustic emissions from a single ear (E–G). The click evoked otoacoustic emissions show that the response differs between ears. The response to a tone burst on the other hand shows evidence of the stimulus periodicity (E–G).

are described first below. The inner hair cells are primarily sensors that ultimately transmit the signal to the brain.

3.3.5.1 Outer Hair Cell

In 1985 it was discovered that outer hair cells are *motile*, that is, able to move (Brownell *et al.* 1985). Since then, the outer hair cells have been the subject of extensive studies. A number of motility mechanisms have been identified. The cell body is able to change its length by up to 5% at high frequencies (above 70 kHz). Some studies have also shown that slower but larger length changes are also possible, but there has also been conflicting evidence on this point. The cuticular plate (labelled in figure 3.6 on page 35) might be able to move relative to the reticular lamina. Hair bundles are able to oscillate. Tip links on the stereocilia seem to be able to actively reposition themselves on the stereocilia.

Much of what we know about the outer hair cells comes from experiments on isolated cells. This allows the conditions of the experiment, such as the ionic concentrations of the surrounding fluid and the cell membrane potential, to be controlled. During these experiments the stereocilia may also be deflected with a glass rod or a water jet, which is an effective input to the hair cell. Adult mammalian outer hair cells suffer considerable damage when they are isolated from the organ of Corti, which means that results using these cells are difficult to interpret. Hair cells from nonmammalian vertebrates, such as frogs and turtles, are hardier and have provided insights into the type of behaviours that might be present in the human cochlea. These experiments show a wide range of responses. Their presence and relative importance in the mammalian cochlea requires further investigation.

Counter intuitively given the vulnerability of the motile outer hair cell response in the cochlea, the motile response of isolated hair cells is not vulnerable. For example, it does not require ATP in the cell², and can survive pressure increases that inflate the cell to a sphere, in addition to the relatively rough treatment required to isolate the cells in the first place. The study of the mechanisms at work to produce this behaviour quickly focuses on the ions of sodium, potassium, and calcium (Na^+ , K^+ , and Ca^{2+}) which control the membrane potential. There are a number of ion channels that selectively pass these ions. The channels may be affected by physical displacements such as transduction channels on the stereocilia, whose opening probability is affected by the tension in tip links between neighbouring stereocilia. They may also be blocked by ions at a molecular level. The behaviour of the gate is described by transduction kinetics equations. A number of important proteins are present in the hair cells. They add structural rigidity where required, and are important for the active response. Myosin motors and actin filaments are present, and prestin appears to be the outer hair cell motor protein in the cell body. Thus, the hair cells in general, and outer hair cells especially, draw their ingenuity from the molecular functions they harness.

²ATP is required by the stria vascularis to maintain the ionic composition of the perilymph and endolymph (see Holley, 1996, page 392).

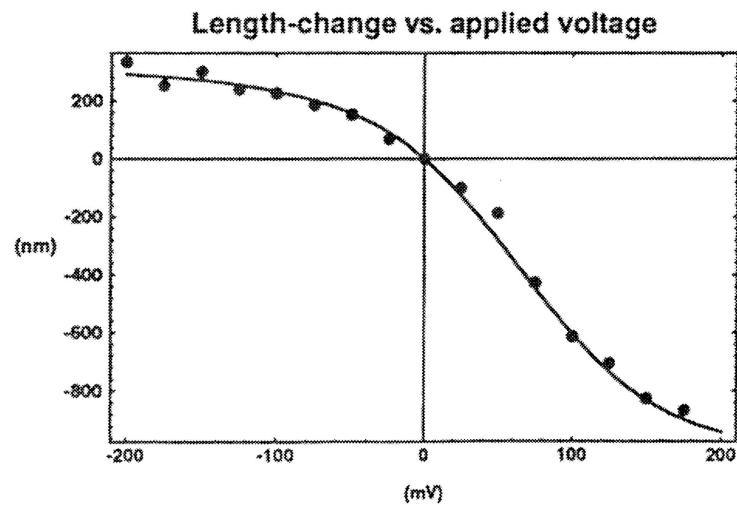


Figure 3.17: Cell length change versus voltage applied to a Guinea pig outer hair cell held in a polished pipette tip. In the setup the pipette is sealed around the cell half way along its length. This leaves the exposed half free to move. (From [Kros et al. 1992](#), figure 7.4b. Work by Dallos and Evans, unpublished.)

Figure 3.17 shows a function of outer hair cell length versus applied voltage for a Guinea pig hair cell held in a pipette. It clearly indicates a variation in length with voltage. Figure 3.18 on page 55 shows an example of adaption. The transducer currents, invoked in a neonatal mouse outer hair cell by a fluid jet on the stereocilia, reduce with time depending on the steady membrane potential (imposed by the experimental setup).

A fuller understanding of the outer hair cell requires further detail on the molecules at work, which is beyond the scope of this review. The reader is instead directed to the reviews: [Kros \(1996\)](#); [Holley \(1996\)](#); [Martin \(2008\)](#); and [Hallworth and Jensen-Smith \(2008\)](#). Some critical points for this work should be noted: the outer hair cells are motile; they are able to respond at high frequencies in excess of 70 kHz; adaption effects on longer time scales exist; and the majority of researchers in the field consider deflection of stereocilia to be the normal mechanical input pathway for the outer hair cells (and hair cells in general). A few researchers have proposed that the outer hair cells may be sensitive to pressure directly (for example, [Bell 2008](#)). While this view is not widely accepted, we will consider its possible implications for the framework at various points to ensure that the framework is able to handle models that require this unusual input pathway.

This discussion shows that we are able to obtain a large amount of data describing the behaviour of the outer hair cells. Nevertheless, our understanding of the role that the outer hair cells play in the active mechanisms in the cochlea is still incomplete. Cochlear models have a contribution to make towards this question by allowing possible mechanisms to be implemented and tested.

3.3.5.2 Inner Hair Cell

The inner hair cells share a number of features with outer hair cells, and hair cells in general. (There are also hair cells in the vestibular system for example.) They use similar ions and have similar ion channels. They also respond to deflection of their stereocilia. However, the inner hair cells are *not* motile. The shape of the cell, and its close proximity to neighbouring cells also suggests this (see figure 3.6 on page 35). Inner hair cells are not able to respond at the high frequencies that the outer hair cell fast motile activity can reach. Instead, their intracellular voltage tracks low frequency signals (less than about 1 kHz) and shows a prominent d.c. component for high frequencies (above 2 kHz). (See [Pickles 1988](#), section 3.D, and figure 3.18 in particular.) The input/output characteristics of inner hair cells that are stimulated at low frequencies show that depolarising voltage changes are larger than hyperpolarising changes (see figure 3.19). Depolarising voltages correspond to rarefaction in sound pressure. Thus, the inner hair cells have a rectifying and low pass filtering action on the signal, to a first approximation.

3.4 Species Dependence

Fitting models of the cochlea to measured behaviour requires a large amount of data. The required data are often not available for a specific animal. Fortunately, the anatomy and physiology of the cochlea is similar among mammals. This means that where data are not available for one species, data from another species can be substituted. However, the data needs to be corrected to account for the known differences between the species; typical examples include the cochlear geometry and frequency map. The fact that mammalian cochleae share strong similarities also allows results obtained on other animals to be used to understand the human cochlea.

The ears of fish, birds, and reptiles are significantly different to mammalian cochleae, and [de Boer \(1991\)](#) section 1.3 provides a useful overview of these differences. Studies in these species are interesting as they show different and simpler approaches to hearing. These species also rely on hair cells to transduce the signal. Amphibians, reptiles, and fish do not have a basilar membrane. They use part of the vestibular system to detect sound. Birds do have a basilar membrane, but it is less specialised than in mammals. Travelling waves have been found in birds, but they have not been found in some other animals. Other mechanisms for frequency sensitivity that have been found where the travelling wave has not been observed include: resonance in the hair cell stereocilia; and electrical resonance based on the ionic permeability of cell membranes ([Crawford and Fettiplace 1981](#); [Ashmore 1983](#); [Lewis and Hudspeth 1983](#)). Hence, data and mechanisms from these species are useful for comparison, but cannot generally be mapped to mammals. The obvious differences also show that different species are striking different balances for the trade-offs discussed in chapter 2.

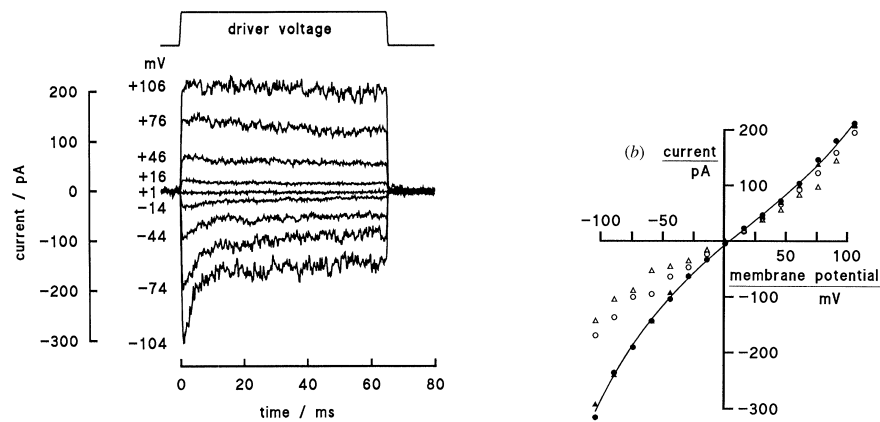


Figure 3.18: Transducer currents in a neonatal mouse outer hair cell. The stereocilia of the cell were stimulated by a fluid jet controlled by the driver voltage indicated in (A). Transducer currents were measured with a patch-clamp technique. This technique allows the membrane potential to be controlled. Part (A) shows unaveraged time responses of the transducer current at various membrane potentials. The response shows adaption to the stimulus for potentials less than -14 mV but no adaption for larger potentials. (B) Current-voltage curves for the cell before and after adaption effects. Filled symbols indicate values before adaption and empty after adaption. Circles indicate values from an experiment where the membrane potentials were systematically incremented and triangles indicate that membrane potentials were systematically decremented. (From [Kros 1996](#), figure 6.3.)




Figure 3.19: The voltage response of a Guinea Pig inner hair cell to 300 Hz tones. Rarefaction pressures are plotted as positive. The graph is asymmetrical, with the response to rarefactions about three times larger than the response to compressions. (From [Russell and Sellick 1983](#), figure 1b.)

3.5 Ethics

Given the vulnerability of the active mechanisms in the cochlea, studies of the physiology of the cochlea in cadavers is of limited use. We are generally interested in the normal operation of healthy ears, and this means that experiments need to be performed on healthy subjects. Therefore, the type of experiments that can be carried out with human subjects is fairly limited, although otoacoustic emissions provide certain opportunities. A range of important psychoacoustic experiments can also be performed, because human subjects can give verbal feedback. These are useful noninvasive techniques; however they inevitably measure the behaviour of the cochlear structures indirectly.

More invasive techniques have been used on animal subjects. Studies of the cochlear input impedance (Lynch III *et al.* 1982) and the response of the chinchilla basilar-membrane to tones (Ruggero *et al.* 1997) will be taken as typical examples of such invasive experiments. In both of these experiments, animals were anaesthetised and the temporal bone opened to expose the cochlea. Openings were also made in the bony wall of the cochlea to accommodate a pressure transducer, and to insert reflective beads on the basilar membrane respectively. The openings were sealed for the work by Lynch *et al.* They were only sealed for some of the experiments by Ruggero *et al.* Precautions to ensure that the condition of the animal is maintained are taken – for example maintaining the body temperature with a heating pad – because the experiments take an extended period of time. Ruggero *et al.* describe monitoring the compound action potential to ensure that the cochleae were in good condition for the duration of the experiment. Once experiments of this kind are completed, the animals are “sacrificed” or killed while still under anaesthetic. Typically this is done by giving an overdose of the anaesthetic. The work by Ruggero *et al.* required 129 animals, and Lynch *et al.* present data from at least 14 and possibly about 25 animals. The experiments would be required to conform to ethical guidelines, and a comment about which body monitors this is often included in the paper. Both of the papers discussed were published in the Journal of the Acoustical Society of America, and therefore they are required to conform with the ethical standards of that publication, which in turn are based on international standards (CIOMS 1985).

The papers above were selected for discussion because they demonstrate the care that is taken in handling the animals, and they have been useful in this particular work. However, the fact remains that a large number of animals have been used in cochlear research, and the question of whether it is right to do this type of research must be addressed. Cochlear models inevitably utilise data from similar experiments, so a position on the ethics of these experiments must be taken.

The three Rs: replacement; reduction; and refinement, are widely used principles for ensuring that humane practices are followed. These were first introduced in Russell and Burch (1959). They refer to designing experiments with the aim of replacing the need for an experimental animal by using an alternative (this is an important role that cochlear modelling aims to fill), reducing the number of animals required to obtain the information, and/or refining the techniques used in the experiment

so that the animal experiences the least possible suffering. Typically when research groups studying the cochlea perform animal experiments they are designed to test a range of things on each animal to ensure that the maximum amount of information is obtained – in itself reducing the need for further experiments.

The principles outlined in [Russell and Burch \(1959\)](#) do not discuss the ethical dilemma of whether it is correct to use animals to further scientific knowledge. The “International Guiding Principles for Biomedical Research Involving Animals” ([CIOMS 1985](#)) states that “Animal experiments should be undertaken only after due consideration of their relevance for human or animal health and the advancement of biological knowledge.” This reflects a utilitarian approach to the question, where the suffering of animal subject is balanced against the expected benefit to human and other animal health. There are examples where action deemed to be ethical in utilitarianism would not be acceptable to most people, so this approach must be used with caution. [Purchase \(2002\)](#) discusses utilitarianism and three other approaches: rights-based ethics; virtue ethics; and deontology or duty-based ethics, and their application to biosciences. He highlights the strengths and weaknesses of each of these positions and notes that while there are other ethical approaches, they are of limited use in bioethics. Most committees set up to assess whether proposed experiments are ethical generally base the decision on utilitarianism, but groups that oppose the use of animals in research generally argue from a rights-based position. [Caplan \(1983\)](#) provides some insight into the arguments around whether animals have intrinsic rights in his discussion of this question.

Above, it was argued that a position on the ethics of animal experiments needs to be taken, when data from animal experiments are used. It is clear that much of what we understand about the ear has been learnt through animal experiments, and it is difficult to see how they could have been fully replaced. Inevitably, a personal position must be taken on this issue when working in this field. Once a (personal) position on the ethics of animal experiments has been established, the question arises of what to do in the case that useful data that does not meet the personal position is found. Not to use the data on the basis that is viewed as obtained through unethical practices requires that the data is sourced elsewhere, which is likely to conflict with the “R” standing for reduction. On the other hand, using the data might be viewed as condoning the original work.

Secondly, in the quote from [CIOMS \(1985\)](#) above it is seen that international best practice only deems that research is justified only if it improves health to some degree. Therefore, experiments that are solely used to improve engineering knowledge are not ethical according to these guidelines. However, where the data exists, and was obtained for valid medical reasons, using the data to advance engineering knowledge for the benefit of humans (and possibly animals) would seem to be a natural way to maximise the value of the experiment.

3.6 Summary

In this chapter the anatomy and physiology of the cochlea has been presented. However, the cochlea cannot be studied in isolation from the rest of the auditory pathway, because its behaviour is so intricately connected to the system. Therefore, an overview of the stages before and after the cochlea was also provided. The structures present in the cochlea were discussed, primarily in the anatomy subsection. Many of these features make regular appearances in biologically inspired models of the cochlea, and their physical parameters such as size, mass, and stiffness are taken directly from anatomical studies where possible. The anatomy also provides some clues about the cochlea's behaviour. We moved on from the anatomy to discuss the physiology of the cochlea. This was tackled by distinguishing between the active and passive behaviour of the cochlea, although it is acknowledged that this useful distinction is hard to define in practice. The passive behaviour was described by highlighting the key characteristics of its response. The active behaviour of the cochlea is less well understood. Rather than discussing any of the models proposed for this behaviour, some of the key types of experiments that demonstrate activity were described. The hair cells are particularly important in the active response. Some key facts about them were described separately. Finally, the ethical concerns surrounding the research that provides us with such valuable information about our sense of hearing were acknowledged and discussed. This chapter has hopefully provided a basis for readers who are unfamiliar with the cochlea to understand the following discussion. It will also be a point of reference when we consider the features that cochlear models might need to include: a cochlear modelling framework should provide mechanisms that allow all the features discussed in this chapter to be simulated.

Chapter 4

Specification – Implications of Cochlear Physiology

This thesis considers the topic of what might be required from a general cochlear modelling framework. Chapter 2 noted the diversity of users of, and uses for, a model. Widely diverse implementations have already been tested. Given the diversity of implementations, we might think that a single cochlear modelling framework cannot possibly be sufficient. However, it is ultimately the physiology of the cochlea that places demands on the framework, and not the implementations that have been used. This chapter, therefore, applies the knowledge of the physiology of the cochlea that was discussed in chapter 3 to determine key features that the framework must be able to support. While the cochlea is undeniably a highly complex organ, it has certain key properties that allow us to place reasonable restrictions on the features which a *cochlear* model must contain. It begins by introducing the concepts of object orientated programming in further detail, and providing an overview of the computational objects that may be selected to model the cochlea. Digressions to discuss the tools available for time domain integration and parallelisation of programs are taken during the course of the chapter. Finally, a summary is provided.

4.1 Objects and Inheritance

The key principle of object orientated programming is to use *objects* encapsulating both the data that describes them, and the methods they can perform. The concept of an object is in itself quite abstract, and various types of objects can be chosen to suit the simulation. Our aim is to find a framework for cochlear modelling that allows people from a wide variety of backgrounds to make a contribution. In this section, we overview the computational objects that were selected. They were chosen to have a close correspondence to physical objects, rather than seeking abstract objects that

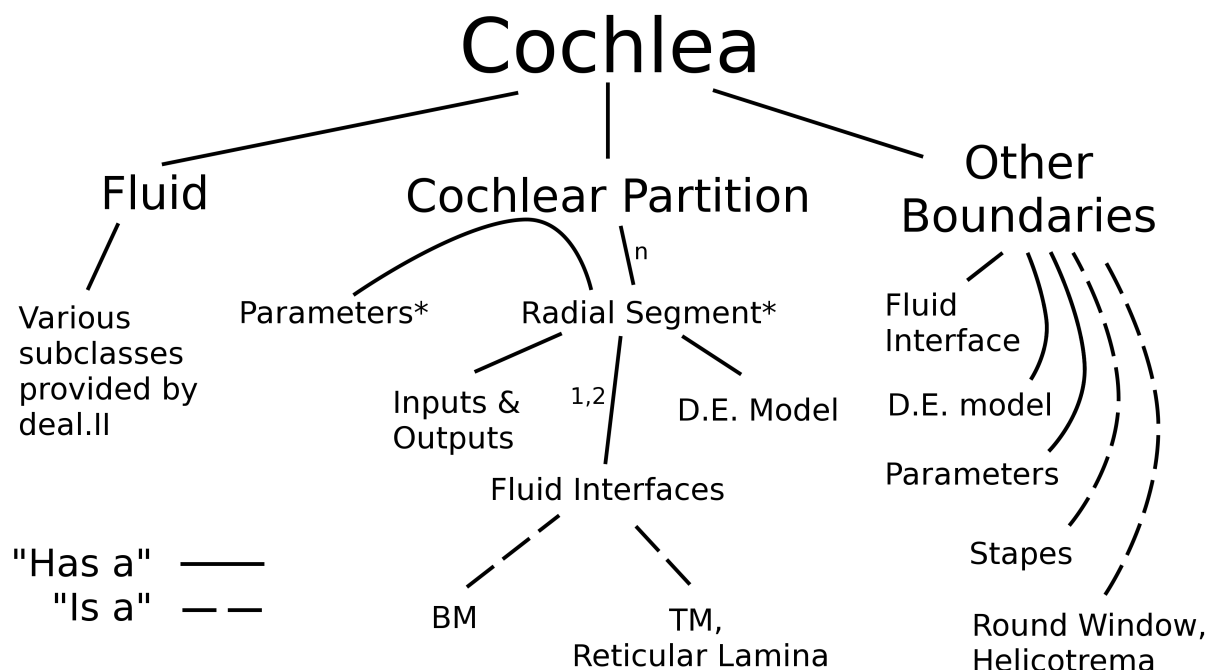


Figure 4.1: An inheritance diagram showing how the cochlea may be broken down into objects. Objects that are expected to be specialised are indicated by either an asterisk (*) after the name, or by examples of the objects set in an “is a” relationship. The numbers 1,2 and the letter n indicate multiplicity: the cochlear partition may have n radial segments; and the radial segment may have one or two fluid interfaces. BM, basilar membrane; TM, tectorial membrane.

might improve computational efficiency at the expense of clarity. This close correspondence is also good practice in the object orientated paradigm (see Bouzeghoub *et al.* 1997, chapter 2).

An object’s form is specified by the *class* to which it belongs. The benefit of this discipline is seen when new classes are created from old either by process of *specialisation* (an “is a ...” relationship, the inverse of which is called *generalisation*), or by *aggregation* (a “has a ...” relationship). *Inheritance* allows classes in a specialisation/generalisation relationship to make use of the data and methods of the more general class, while selectively redefining some methods (via *polymorphism*). An object made up by the aggregation of other objects leverages the behaviours that they offer to perform its function. (See the reference above for more detail about these terms.)

Figure 4.1 shows how the cochlea is broken down into a series of objects for simulation. The cochlea is first divided into its major constituent parts: the fluid; cochlear partition; and other boundaries (of which the stapes, and the round window or helicotrema are examples). These objects in turn comprise of other parts. We state the major parts now, and justify them in the discussion that follows.

The behaviour of the cochlear partition is hard to express as a single object, but the response of a short section to inputs can be specified. Therefore, by using an aggregation of simpler radial segment objects, we avoid this complexity. A radial segment requires physiological data, and a parameters

object is provided to the segments by the cochlear partition object as required. The radial segments in turn make use of objects, including some to handle its interface with the fluid (either one or two depending on the model) and differential equations that express its dynamics. The other boundaries of the fluid region, which include the stapes, and the round window or helicotrema, also require similar sub-objects.

The sub-objects required to implement the fluid are provided by the deal.II library, and will not be discussed at this stage. Instead, the discussion concerning the cochlear fluid will focus on the suitability of the finite element method to represent this object.

4.2 The Cochlear Fluid

The cochlear partition is the focus of current research on the cochlea. However, the fluid in the cochlea plays the key role of transmitting the sound input to the cochlear partition, and coupling the various parts of the partition. Therefore, while our interest lies in the behaviour of the cochlear partition, the fluid region serves as a natural foundation for any cochlear model.

To date, the mechanical function of the fluid has been considered in cochlear models, and its function of transporting ionic species to and from the cochlear partition has been largely neglected. A recent model ([Ramamoorthy *et al.* 2007](#)) includes a one-dimensional (1D) model of the macroscopic current behaviour, so we anticipate that this will become more important in the future. The majority of this discussion will focus on the fluid dynamics, however.

The starting point for computation fluid dynamics (CFD) techniques is a set of partial differential equations that describe the behaviour of the fluid in question. Max Viergever studied the widely used assumptions that the cochlear fluid is incompressible, inviscid, and linear in his PhD thesis ([Viergever 1980](#)). He started with the general equations for a Newtonian fluid, and using dimensional analysis, quantified the conditions under which each assumption is valid. His logical reduction of the original equations to a more manageable form is still important to understand today, as it lends support to these widely used assumptions. Therefore, we begin by retracing his argument in summarised form, before updating the argument with more recent literature. We also address fluid dimension and mesh moving, finally noting that the finite element method is adequate to handle the full spectrum of fluid models that might be required.

4.2.1 Incompressibility – Viergever’s Contribution¹

The general governing equations for a Newtonian fluid are:

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{v}) = 0, \quad (4.1)$$

$$\rho \frac{\partial \mathbf{v}}{\partial t} + \rho (\mathbf{v} \cdot \nabla) \mathbf{v} + \nabla p - \mu \nabla^2 \mathbf{v} - \left(\zeta + \frac{1}{3} \mu \right) \nabla \nabla \cdot \mathbf{v} = \mathbf{0}, \quad (4.2)$$

where ρ is the fluid density, \mathbf{v} is the velocity, p is the pressure, μ and ζ are the coefficients of viscosity, and ∇ is the Laplace operator. To study the effect of compressibility, Viergever assumes:

$$c^2 = \frac{dp}{d\rho}, \quad (4.3)$$

where c is the speed of sound in the fluid. Using (4.1) and (4.3), he obtains:

$$\frac{1}{c^2} \frac{\partial p}{\partial t} + \frac{1}{c^2} (\mathbf{v} \cdot \nabla p) + \rho \nabla \cdot \mathbf{v} = 0. \quad (4.4)$$

He proceeds with the dimensional analysis by selecting a characteristic time and length (T_c and L_c respectively) where these characterise the order of magnitude required in each quantity for a significant change in the velocity. He also defines the characteristic amplitude A_c and angular frequency ω of the oscillations. Using these, he shows that the following relationships must hold for compressibility to be negligible:

$$\omega^2 \ll \frac{c^2}{L_c^2}, \quad (4.5)$$

$$\omega^2 \ll \frac{c^2}{A_c L_c}. \quad (4.6)$$

Viergever assumes 1500 ms^{-1} for c and 7 mm for L_c (based on the largest diameter of the human cochlea). He also notes that $A_c \ll L_c$. Viergever specifies the audio range² as $10^2 < \omega < 10^5 \text{ Hz}$. He concludes that equations (4.5) and (4.6) are satisfied for all except the highest frequencies in the audio range.

¹In this series of headings, I am not suggesting that Max Viergever was the first to use these ideas. In fact, they date to the start of the field. His logical argument for these ideas is the contribution that I refer to.

²Viergever stated this range in cycles per second, despite defining that $\omega = 2\pi f$ earlier.

4.2.2 Linearity – Viergever’s Contribution

Having considered the incompressibility assumption, Viergever continued to assess the linearity assumption using the Navier-Stokes equations:

$$\nabla \cdot \mathbf{v} = 0, \quad (4.7)$$

$$\rho \frac{\partial \mathbf{v}}{\partial t} + \rho (\mathbf{v} \cdot \nabla) \mathbf{v} + \nabla p - \mu \nabla^2 \mathbf{v} = \mathbf{0}. \quad (4.8)$$

Equation (4.7) is the equivalent of (4.1) for an incompressible fluid, and (4.8) is the Navier-Stokes equation, which replaces (4.2). Viergever compared the magnitude of the linear convection term $\rho \partial \mathbf{v} / \partial t$ to the advective term $\rho (\mathbf{v} \cdot \nabla) \mathbf{v}$. Once again he used dimensional analysis, and determined that the fluid can be treated as linear if:

$$A_c \ll L_c. \quad (4.9)$$

He suggests a value of 1 mm for L_c , based on either the cross-section of the scala or minimum wavelength of displacement waves on the basilar membrane. For A_c , he suggests values of 10 μm and 0.1 μm for input sound pressure levels (SPL) of 140 dB and 100 dB (re 20 μPa) respectively. He concludes that the fluid behaves linearly, because relationship (4.9) holds even for loud sound pressure levels.

4.2.3 Viscosity – Viergever’s Contribution

Finally, Viergever compares the magnitude of inertial forces ($\rho \partial \mathbf{v} / \partial t$) to viscous forces ($\mu \nabla^2 \mathbf{v}$) using the Reynolds number for the situation that he derives using dimensional analysis:

$$Re = \frac{\text{inertial forces}}{\text{viscous forces}} \approx \frac{\rho \omega L_c^2}{\mu}. \quad (4.10)$$

Using values for the density and viscosity of 1 $\text{mg} \cdot \text{mm}^{-3}$ and 1 $\text{mg} \cdot \text{mm}^{-1} \cdot \text{s}^{-1}$ respectively, and 1 mm for L_c as before, he infers that the Reynolds number is of similar order of magnitude to ω . Therefore, he concludes that away from boundary walls, the fluid may be treated as inviscid in the audio range. Near solid walls, a viscous boundary layer is always present. Viergever provides an estimate for the thickness of the boundary layer:

$$\delta \approx \left(\frac{\mu}{\rho \omega} \right)^{\frac{1}{2}}, \quad (4.11)$$

where δ is the boundary layer thickness. When ω is sufficiently large, $\delta \ll L_c$, and the viscous boundary layer is negligible (Viergever also gives formulae that can be used to correct for its presence). However, for frequencies below the audio range this may not be satisfied.

4.2.4 Viergever's Contribution – Implications

The work that Viergever presented, which has been summarised above, lends support for three important assumptions about the fluid in the cochlea. The significance of the assumptions is that they allow the vector problem of equations (4.1–4.2) or (4.7–4.8) to be replaced by a scalar problem. This reduces the computational expense of solving the fluid domain problem considerably. Two approaches to obtain a scalar equation are possible. They both start from the following equations:

$$\nabla \cdot \mathbf{v} = 0, \quad (4.12)$$

$$\rho \frac{\partial \mathbf{v}}{\partial t} + \nabla p = \mathbf{0}, \quad (4.13)$$

where the linearity of the fluid, and its inviscid nature have been used to simplify (4.8) to obtain (4.13).

Firstly, note that (4.12) implies $\nabla \cdot \mathbf{a} = 0$, where $\mathbf{a} = \partial \mathbf{v} / \partial t$ is the acceleration of the fluid, so:

$$\nabla^2 p = 0, \quad (4.14)$$

$$\mathbf{a} = -\frac{1}{\rho} \nabla p, \quad (4.15)$$

where we have assumed that ρ does not vary spatially for simplicity.

The alternative, which is common in the field of fluid dynamics (Douglas *et al.* 2001), is to propose a velocity potential field, ϕ , such that $\mathbf{v} = -\nabla \phi$. Using this definition, it can be shown (Lyon and Mead 1988b) that:

$$\nabla^2 \phi = 0, \quad (4.16)$$

$$p = \rho \frac{\partial \phi}{\partial t}. \quad (4.17)$$

The differential equations (4.14) and (4.16) can be solved using similar techniques. Which equations to use in a given situation is to some extent personal preference, but they require different methods for postprocessing the data. (4.14) provides the pressure field as part of the solution. To obtain the velocity field, the fluid acceleration must be calculated and then integrated with respect to time. (4.16) does not provide the pressure, which must be calculated using a temporal derivative. On the other hand extracting the velocity field is less computationally intensive, because it is the spatial derivative of the potential field.

In both (4.14) and (4.16) we need to specify either the value of the scalar field (p or ϕ) or its *normal* gradient for any point on the boundary of the fluid domain. We are not able to specify any value for the tangential components. This fact becomes apparent when the equations are converted to boundary value problems. Although this is a known result for scalar field formulations, the underlying

reason why we are unable to specify the tangential components is because of the assumptions we have made: the inviscid fluid implies that the boundary is unable to exert a shear force on the fluid, and hence the tangential component cannot be constrained. This assumption also happens to allow us to simplify the vector problem to a scalar problem.

4.2.5 Viergever's Contribution – Reassessment

The analysis that Viergever used to justify the assumption that the cochlear fluid is inviscid, incompressible, and linear was presented in 1980. The quality of predictions that it makes concerning the behaviour of the fluid is dependent on the quality of the estimates for the characteristic dimensions (L_c , A_c , c , μ , and ρ) of the flow pattern present in the cochlea. 30 years further on, the values he suggested are still good estimates.

Robles and Ruggero state a value of 1550 ms^{-1} for the speed of acoustic pressure waves in the cochlea, c , (Robles and Ruggero 2001). They also provide estimates of the wavelength of displacement waves in various species (one value that could be used for L_c). All these estimates lie in the range 0.5 to 1.8 mm (see their table 4). Finally, their figure 3 shows a plot of basilar membrane displacement (A_c) against sound pressure level for four species. The sound pressure levels extend to 110 dB, and all the resulting displacements are less than 30 nm. Therefore, this paper suggests that the values Viergever selected were appropriate, although they were slightly pessimistic in the case of A_c .

Puria and Allen use the following parameters in Puria and Allen (1991): $\rho = 1 \text{ mg}\cdot\text{mm}^{-3}$; $\mu = 2 \text{ mg}\cdot\text{mm}^{-1}\cdot\text{s}^{-1}$; $c = 1500 \text{ m}\cdot\text{s}^{-1}$. The value that they assume for the viscosity is twice that used by Viergever. They based this value on the viscosity of sea water (personal communication). This is also the same figure that Von Békésy found (Peterson and Bogert 1950). Viergever cites data more recent than Von Békésy's study to justify at his estimate for viscosity. However, a doubling in the value used for the viscosity would only imply that viscous effects become important at slightly higher frequencies.

Finally, we consider the alternative method of estimating L_c based on the dimensions of the cochlear fluid spaces. Thorne *et al.* provide data for the fluid space dimensions from six species (Thorne *et al.* 1999). The human data they present shows that the scala vestibuli and the scala tympani both have cross-sectional areas of approximately 1 mm^2 . Therefore, the value for L_c that Viergever adopts (1 mm) is appropriate for humans. Five other species are listed: the guinea pig; bat; rat; gerbil; and mouse. A cross-sectional area 0.1 mm^2 is more typical of the fluid spaces in these species, suggesting that an L_c of 0.3 mm ($\approx \sqrt{0.1}$) would be more representative for these species. Using this value of L_c for the animals listed would have negligible effect on the onset of nonlinear behaviour in the fluid, because A_c is much smaller still, but it would once again have a slight effect on the onset of viscous effects.

The assumption that the cochlear fluid is inviscid can be problematic in time domain cochlear models. As Viergever shows, this assumption is only valid for frequencies above a certain limit. [Matthews \(1980\)](#) pages 70–75 and [Neely \(1978\)](#) page 56 separately noticed the possibility of low frequency resonances in the inviscid cochlear models that they were using. The problematic frequencies can simply be avoided in frequency domain models, but for time domain models this is not possible. Allen and Puria used a transmission line model that included viscosity to study the input impedance of the cochlea ([Puria and Allen 1991](#)). They found that including viscosity and the tapering of the cross-sectional area of the fluid region reduced these standing waves significantly. They also noted that a nonphysiological cochlear map aggravated the problem, and offered a cochlear map that virtually eliminates them. (However, this cochlea map does not match the best available physiological data.) Alternatively, the helicotrema boundary condition is sometimes used to avoid these resonances in time domain models that use an inviscid fluid model.

[Puria and Allen \(1991\)](#) considered compressibility of the fluid. They found that the effects of compressibility were largest near the base of the cochlea, where they state that that fluid compliance (due to compressibility) is less than 6% of the basilar membrane admittance. After assessing its effects, they neglected it, as it had little impact on their results (they state less than 0.5 dB near 20 kHz). However, [Ramamoorthy *et al.*](#) recently found that including compressibility in a frequency domain model improved the agreement between the frequency dispersion in the biological cochlea and in the model, after the best place. They suggest that this could be because a different mechanism dominates in this region ([Ramamoorthy *et al.* 2010](#)).

The picture that emerges from both Viergever’s dimensional analysis approach and this discussion is that the adequacy of the assumptions cannot be determined once and for all. Rather, whether to use these assumptions or not is a compromise that must be made, which depends on the particular model. The assumptions inevitably introduce errors into the model, but reduce the computational effort required for the model considerably. The errors may be negligible for the frequency range of interest when compared to other errors in the model; for example, those due to uncertainty in the parameters used for the model. Currently they seem to be adequate, but as better data and more computing power becomes available, it seems fair to expect that previously acceptable assumptions will be repudiated. Furthermore, the ability to include these effects to test their importance is also useful. [Puria and Allen \(1991\)](#) and [Ramamoorthy *et al.* \(2010\)](#) are examples of this use. Therefore, while the majority of cochlear models to date assume a linear, inviscid, and incompressible fluid, the flexible cochlear modelling framework developed here should provide for models that do not.

4.2.6 Spatial Dimension

In the sections above, we have discussed the partial differential equations that govern the behaviour of the fluid in the cochlea. We now consider what spatial dimension should be used to model the

geometry of the fluid region. One obvious starting point would be to consider a 3D fluid region, because that is what is actually present in the cochlea. On the other hand, it is interesting that the sensory cells in the organ of Corti (the inner hair cells) are arranged in a single line running from the base of cochlea to its apex. The fact that the output stage is 1D seems to suggest that even a 1D model might be adequate. In fact, 1D, 2D, and 3D models of the cochlea have all been used in various studies. [Diependaal *et al.* \(1987\)](#) is an example of an early time domain model in 1D. [Allen and Sondhi \(1979\)](#), [Matthews \(1980\)](#), [Diependaal and Viergever \(1989\)](#), [Kim and Xin \(2005\)](#), and [Liu and Neely \(2010\)](#) are examples in 2D. Some 3D models include [Steele and Taber \(1979b\)](#), [Kolston and Ashmore \(1996\)](#), [Lim and Steele \(2002\)](#), and [Givelberg and Bunn \(2003\)](#).

To date, 3D models have not been widely used in comparison to 1D and 2D models. A major reason is the increased computational effort these models require. As an example, compare the first finite difference 2D cochlear model presented in [Neely \(1981a\)](#) to the 3D finite element model in [Kolston and Ashmore \(1996\)](#). In the 2D model, 256 points were used in the length direction and 8 points in the height direction, for a total of 2048 mesh points. The smaller of the two meshes used in the 3D model required $60 \times 20 \times 30$ nodes, for a total of 36000 nodes. Despite having about four times fewer nodes in the length direction, and only about twice as many nodes in the height direction, the model still requires about 20 times more nodes in total. This increase in the number of data points required to simulate the cochlea has a direct impact on the amount of RAM and disk space required to run the simulation. Furthermore, the structure of finite element and finite difference matrices are different for 2D and 3D models. The algorithms used to manipulate data such as matrices in the model can scale as n^4 in 2D, but n^7 in 3D, where n is the number of nodes. This is because in a 2D model the coupling between nodes must lie in a plane, whereas for a 3D model the coupling lies in a volume, so each node has more neighbours. Therefore, even for similar sized matrices, operations to construct and solve the 3D model are more expensive than for a 2D model. All of these factors contribute to the increased computational expense of 3D models (see [Strang \(2007\)](#) chapter 7, and appendix C in this thesis).

When selecting to use 256 points in the model described above, Neely also tested 100, 200 and 400 divisions and found that his results were not sensitive to this change. It is interesting to note that in a recent paper by Neely ([Liu and Neely 2010](#)) the number of divisions had increased to 700. It seems natural that it should increase as improved computing facilities allow researchers to demand greater accuracy. (The accuracy of the numerical solution to most partial differential equations improves for an increasingly fine spatial mesh. In the case of the finite element method, as well as other methods, the discrete solution should converge a continuous solution as the mesh is refined.) This raises the question of whether an upper bound can be placed on the number of divisions that might be used in a cochlear model. The number of inner hair cells along the length of the cochlea (about 3000) seems to place a natural limit on the number of divisions that would be useful. This will be picked up again in section 4.4 below.

If increasing the dimension of a model from 2D to 3D increases the computational expense of solving the model, then naturally reducing the dimension from 2D to 1D reduces it. However, there is also a drawback in reducing the dimension of a cochlear model. Viergever also investigated the adequacy of 1D models compared to 2D and 3D models. He used both theoretical arguments (see [Viergever 1980](#), appendix 2c) and numerical results (see [Viergever 1980](#), chapter 5.2.E). He concluded that while a 1D model gives a reasonable qualitative match to a 2D model, the quantitative agreement between the models is poor. He also found that the 2D model which he studied was a fair approximation to a 3D model: the solution to the 2D model was the average of 3D model over the omitted dimension. Nevertheless, the 1D model is still useful, both to prove general results (for example, [de Boer and MacKay 1980](#)) as well as for numerical calculations (for example, [Ku et al. 2008](#)).

One of the advantages of treating the cochlear fluid as linear, inviscid, and incompressible is that simulations with 2D or 3D accuracy in the fluid model which have a computational cost similar to a 1D model can be constructed. One way to achieve this is to use Green's functions as in [Allen \(1977\)](#), [Matthews \(1980\)](#), and [Diependaal and Viergever \(1989\)](#). An approach using the finite element method will be presented in chapter 5.

This section has shown that the choice of spatial dimension is an important trade-off when designing a cochlear model. Higher spatial dimensions have improved accuracy, but come at a computational cost. Currently, a 2D model seems to be a happy medium, but a cochlear modelling framework should also offer the ability to implement 1D and 3D models.

4.2.7 Fluid Geometry

The choices available in terms of a set of partial differential equations, and spatial dimension for the fluid model have been discussed. The specific geometry of the fluid region must also be considered. Section 3.2 introduced the scala vestibuli, scala media, and scala tympani, and noted that the scala media can be lumped with the scala vestibuli from a mechanical point of view. It is common to assume that the scala vestibuli and scala tympani are symmetrical. Using this assumption, [Peterson and Bogert \(1950\)](#) showed that only the pressure difference between the two scalae is responsible for the cochlear travelling wave. The common mode pressure has no effect on the basilar membrane motion in this case, so it can be neglected. The differential mode pressure is symmetrical about the basilar membrane, $p(x, y) - p(x, -y) = 2p(x, y)$. This implies that only one of the scalae needs to be modelled – halving the size of the simulation.

While this reduction in problem size is attractive, there are cases where both scalae should be included in the simulation. The first is if the symmetry condition on the scalae is not met. [Mountain et al. \(2003\)](#) provide data for the cross-sectional area of the cat cochlea. The ratio of cross-sectional areas in the scala vestibuli plus scala media to the area in the scala tympani is as high as 1.86 in the “upper base” of the cochlea, and 107.5 in the apex (this is an extreme value compared to other

species). Therefore, while the assumption that the scalae are symmetrical is widely used, more accurate fluid geometries that do not require this simplification should ultimately give an improved match to the data, if accuracy is essential. Secondly, neglecting one of the scalae in simulation assumes that the cochlear partition is incompressible. Some models require the cochlear partition to compress, in which case this assumption is not justified (Flax 2008). Finally, some authors argue that the common mode pressure plays a role in exciting the outer hair cells directly (Bell 2005, 2008; and Braun 2009). Overall, the modelling framework should be able to cater to models that simplify the geometry to a single scala, as well as those that do not.

A second common simplification is to approximate the scala shape by a rectangle (or rectangular prism in 3D). This approximation is particularly useful when the Liouville-Green approach (Steele and Taber 1979a) or the Green's function method (Allen 1977) are used because the geometry simplifies the integrals that need to be solved. (A recent paper (Lüling *et al.* 2010) introduces conformal mappings that appear to share a close relationship to the Green's function method, and this approach allows the geometry to be specified more freely.) Once again, the literature suggests that a more complex geometry may improve the accuracy of a simulation. In the discussion on viscosity above, it was noted that including scalae tapering in a model can reduce apical reflections (Puria and Allen 1991). Cai *et al.* (2005) investigated the effects of the coiling of the scalae in the cochlea. They note significant effects, particularly in the apical region where the curvature is greatest.

4.2.8 Fluid-Structure Interaction and Mesh Moving

Fluid-structure interaction is a hot topic in the computational fluid dynamics literature at the present moment. It is also important in cochlear modelling because the cochlear partition, stapes, and round window all lie in direct contact with the cochlear fluid, and move in response to the fluid's motion. The coupling occurs in both directions, however, because these structures also affect the motion of the fluid. Formally, the cochlea shows "strong fluid-structure interaction", which has important repercussions on the structure of the framework that will be presented. However, a discussion of these repercussions will be delayed until chapter 5. For the rest of this section, the framework will be specified as if the fluid-structure interaction was not problematic.

In general, a deflection in boundary of a fluid region implies that the shape of fluid region has also undergone a change. When the change in shape is significant relative to the original size of the region, the change must be accounted for in the simulation. In a cochlear model, however, the maximum displacement of the various structures is much smaller than the size of the fluid region. The diameter of the scalae in humans is roughly 1 mm, and the amplitude of basilar membrane vibrations are less than 100 nm even for inputs of 100 dB SPL (see section 4.2.5). Stenfelt *et al.* (2004a) provide data for the volume displacement of the oval and round windows. They find that the volume displacement is similar for both windows – roughly $3 \times 10^4 \mu\text{m}^{-3}$. For comparison,

they also state an estimate for the volume of the fluid in the cochlea: 98.1 mm^{-3} . Therefore, the maximum displacements of these structures is negligible (and will remain negligible) in relation to the dimensions of the cochlea.

4.2.9 Conclusion

The cochlear fluid plays a key role in coupling together the various structures in the cochlea. Therefore, a physiologically accurate cochlear model requires an object to represent it. As this section has shown, the requirements placed on this object can vary considerably depending on the approximations deemed acceptable for the partial differential equations describing the fluid's behaviour, and the spatial dimension adopted for the model. The arguments presented in section 2.2 showed that some models aim to capture the essential characteristics of cochlear physiology at the expense of accuracy in the fit to physiological data, whereas others aim for the best fit possible at the expense of complexity in the model. In setting up this framework, both of these trade-offs (and the other trade-offs discussed) are viewed as valid, and care is taken to cater for both where possible. The finite element method has been successfully applied to cochlear modelling in the past. Furthermore, the diversity of the problems that it has been applied to in other fields shows that it is able to handle the variations discussed in this section, and more. Therefore, the finite element method is a suitable tool to use for the fluid object.

4.3 State Space and Time Integration

Having discussed the cochlear fluid, we move on to discuss the structures around the fluid, and in particular, the cochlear partition. Whereas a Laplacian fluid model (see equations (4.14) and (4.16)) contains no time dependent terms, time dependence is a key feature of these structures. State space is a method with sufficient flexibility to express the types of models required for the cochlear partition. In addition to the the state space representation, we also require a time stepping algorithm such as the Runge-Kutta 4th order method. Therefore, we take a digression to discuss these tools that will form an important part of the framework.

4.3.1 Definitions

4.3.1.1 State Space

MacFarlane provides a good overview and history of state space systems in chapter six ("State Models") of his book [MacFarlane \(1970\)](#). He notes that the concept of a state was introduced by

Poincaré in the 1890s. Poincaré extended work by Cauchy and Moigno on differential equations. MacFarlane provides the following definition of a system's state:

A state of a dynamical system is any complete independent set of system model variables whose knowledge enables one to predict completely the future behaviour of the system.

For many models, the state of a system can be represented by a vector of variables whose behaviour is governed by first order differential equations. When high order differential equations are studied, auxiliary variables can be introduced to separate it into a set of first order differential equations. The success of this method depends largely on the fact that basic energy storage devices are governed by first order differential equations, so a correspondence can be set up between the state of a system and its energy storage elements. However, we are not restricted to selecting the energy storage elements as the state. A suitable mapping may be used to redefine the state of the system, provided that the resulting state is complete and independent.

Based on a state $\mathbf{X} = \begin{bmatrix} X_1 & X_2 & \cdots & X_{n_s} \end{bmatrix}^T$, and set of external forcing values $\mathbf{U} = \begin{bmatrix} U_1 & U_2 & \cdots & U_r \end{bmatrix}^T$, the set of state update equations may be expressed as:

$$\frac{dX_i}{dt} = f_i(X_1, X_2, \dots, X_{n_s}; U_1, U_2, \dots, U_r; t), \quad i = 1, 2, \dots, n_s. \quad (4.18)$$

This equation governs the dynamics of the state of the system, which directly or indirectly represent the energy storage of the system, as we noted. The behaviour of (a) specific part(s) of the model may be monitored using a (set of) output equation(s) such as:

$$y_i = g_i(X_1, X_2, \dots, X_{n_s}; U_1, U_2, \dots, U_r; t), \quad i = 1, 2, \dots, n_o, \quad (4.19)$$

where n_o is the number of output variables of interest. In the special case that an output of interest is also in the system state set, the specific output equation would simply select that desired variable from the set³. Note that the state update equations are a set of *differential equations*, whereas the output equations contains no derivatives.

4.3.1.2 The Runge-Kutta 4th Order Algorithm

Equations (4.18) and (4.19) provide us with a flexible *representation* of a dynamical system. In order to use this representation in a cochlear model, we require a method to solve the differential equation (4.18). The Runge-Kutta 4th order algorithm provides us with one alternative to do so. It simulates the response of the system at a set of times by calculating the state one time step (Δt) later, based on

³When it will not cause undue confusion, a variable in the system state set will be referred to simply as a state.

the values at the current time. The algorithm works by computing four estimates of the derivative dX_i/dt over the period Δt , and combining them using a weighted average to provide an improved estimate for the gradient of the system state. Finally, this improved estimate is used to calculate the states for the next time step. (See [Butcher and Wanner 1996](#) for an interesting review of the history of this method.) Formally, the algorithm can be represented as:

$$\dot{\mathbf{X}}_1 = \mathbf{f}(\mathbf{X}(t), \mathbf{U}(t), t), \quad (4.20)$$

$$\mathbf{X}_1 = \mathbf{X}(t) + \frac{1}{2}\Delta t \dot{\mathbf{X}}_1, \quad (4.21)$$

$$\dot{\mathbf{X}}_2 = \mathbf{f}\left(\mathbf{X}_1, \mathbf{U}\left(t + \frac{1}{2}\Delta t\right), t + \frac{1}{2}\Delta t\right), \quad (4.22)$$

$$\mathbf{X}_2 = \mathbf{X}_1 + \frac{1}{2}\Delta t \dot{\mathbf{X}}_2, \quad (4.23)$$

$$\dot{\mathbf{X}}_3 = \mathbf{f}\left(\mathbf{X}_2, \mathbf{U}\left(t + \frac{1}{2}\Delta t\right), t + \frac{1}{2}\Delta t\right), \quad (4.24)$$

$$\mathbf{X}_3 = \mathbf{X}_2 + \Delta t \dot{\mathbf{X}}_3, \quad (4.25)$$

$$\dot{\mathbf{X}}_4 = \mathbf{f}(\mathbf{X}_3, \mathbf{U}(t + \Delta t), t + \Delta t), \quad (4.26)$$

$$\dot{\mathbf{X}}_{avg} = \frac{1}{6}\dot{\mathbf{X}}_1 + \frac{1}{3}\dot{\mathbf{X}}_2 + \frac{1}{3}\dot{\mathbf{X}}_3 + \frac{1}{6}\dot{\mathbf{X}}_4, \quad (4.27)$$

$$\mathbf{X}(t + \Delta t) = \mathbf{X}(t) + \Delta t \dot{\mathbf{X}}_{avg}, \quad (4.28)$$

where $\dot{\mathbf{X}}_i$ and \mathbf{X}_i are intermediate estimates of the state gradient and state respectively, Δt is the time step size, and the notation $\mathbf{X}(t)$ and $\mathbf{U}(t)$ imply the value of the state and input at time t .

The Runge-Kutta 4th order method is one of a family of time stepping methods. It is popular in cochlear modelling because it is cheap to compute and, as the name implies, *fourth-order convergent*, which means that the error in the estimate for $\mathbf{X}(t + \Delta t)$ decreases as $\Delta t^{(4+1)}$ when Δt is reduced, assuming $\mathbf{X}(t)$ is exact. Diependaal *et al.* considered various alternatives, and determined that the Runge-Kutta 4th order method provided the best trade-off between speed and accuracy ([Diependaal et al. 1987](#)). When this method is applied to an object in isolation, the state represents variables internal to the object, while the inputs are affected by factors external to the object. Note that this method requires estimates for $\mathbf{U}(t)$ at the current time t , the final time $t + \Delta t$, and twice at $t + \Delta t/2$. If there is a feedback path between the state \mathbf{X} and \mathbf{U} such that the value of \mathbf{U} is significantly affected by changes in \mathbf{X} over the course of a single time step, then the values of $\mathbf{U}(t + \Delta t)$ required for the second and third estimate are different. This is discussed further in chapter 5, because the fluid forms an important feedback path.

4.3.2 Formal Limitations on these Methods

4.3.2.1 Existence and Uniqueness of the Solution

The methods described above are suitable for a wide range of dynamical models, but some factors need to be considered. Firstly, equation (4.18) is not guaranteed to have a solution for an arbitrary set of functions $\{f_i\}$ and specific initial conditions \mathbf{X}_0 . That is, the solution may not *exist*. During a simulation, if a solution does not exist for a given \mathbf{X}_0 the algorithm will either fail (if we are lucky) or may produce an incorrect result without giving any obvious indication. Furthermore, even if a solution starting at (or passing through) \mathbf{X}_0 does exist, it need not be the only valid solution; that is, it need not be *unique*. If we attempt to find a solution to a problem that does not have a unique solution using a simulation, we only find one of the possible solutions. We would have not completely characterised the behaviour that might arise because the remainder of the solutions have not been found. Once again, we may not have an indication that this is the case. [MacFarlane \(1970\)](#) section 6.1 addresses this question, and provides a proof that if the state update functions are *continuous* and satisfy a *Lipschitz condition* at a specified point, then a unique solution exists at that point and in its vicinity. Specifically, continuity is used to show existence, and the Lipschitz condition, together with continuity, to show uniqueness. These conditions are stricter than generally required, as they guarantee that the solution also has continuous derivatives. They can be relaxed slightly to allow $\{f_i\}$ to be discontinuous in time (see [Coddington and Levinson 1955](#), pg 42). Even these conditions are too strict for some applications of differential equations (notably optimal control theory). Therefore, alternative definitions of what we mean by a *solution* to a differential equation may be used to extend the idea to $\{f_i\}$ that are discontinuous with respect to the state of the system as well as time ([Hájek 1979](#)).

This discussion highlights the fact that the state space representation is very flexible. It also serves as a caution that the representation provides little in the way of safeguards that a problem is well-posed. When this framework is used to compare successful models already described in the literature, this is of less concern, because the conditions are likely to be met. However, it is of concern when new models are created using the framework. As we noted in section 2.3.3, flexibility and power are given precedence in this cochlear modelling framework, so no attempt is made to exclude models that may be poorly posed. Instead, we make a point to discuss the danger. Many cochlear models satisfy the continuity and Lipschitz conditions (which do not automatically exclude nonlinear functions), in which case, the existence and uniqueness of the solution is guaranteed.

4.3.2.2 Causality

Secondly, the system should be *physically realisable* or *causal*. [Stremmer \(1990\)](#) defines it in the following manner:

A physically realisable system cannot have an output response before an arbitrary input function is applied. Stated another way, the output of a physical system at $t = t_0$, namely, $g(t_0)$, must depend only on values of the input $f(t)$ for $t \leq t_0$. A system having this property is called *physically realisable or causal*. Any system not meeting this requirement is said to be nonrealisable or noncausal.

It might seem that placing a causality requirement on the cochlear modelling framework is not a significant limitation, because our ambition is to simulate a real system that must be causal by definition. However, noncausal systems are useful both as theoretical tools and in practice. Audio compression algorithms such as mp3 make use of future samples to obtain a better encoding (Brandenburg 1999). They can do this because the samples are available ahead of time in the recording. Noncausal systems can also be studied in order to find an optimal solution as an intermediate stage before a physically realisable approximation is found. In the language of section 2.1.2, a noncausal system can be a perfectly valid abstract model, and may even be our ideal solution, but the implemented model can at best approximate it. An example from control theory is as follows: A design process such as using a root-locus diagram may suggest that a proportional-differential (PD) controller is optimal for a given application. The derivative term in this controller requires an ideal differentiator, which cannot be realised in practice, either in hardware or software, because it is noncausal. Instead, the differentiator is augmented with a nondominant pole that allows the controller to be realised (at the cost of the realisable controller slightly underperforming the theoretical design). This discussion highlights the fact that methods which allow the properties of noncausal systems to be studied are useful despite the fact that all physical systems are causal.

It is worth asking whether this modelling framework could allow noncausal objects to be modelled directly. The state space representation and Runge-Kutta time stepping algorithm currently limits the simulation to causal models. The causality requirement could possibly be relaxed by selecting a different set of tools for the time integration. However, noncausal blocks by definition require some future inputs up front. This in itself places new restrictions on the source of input signals and the interaction between the objects, which in turn limits the generality of the blocks that this framework attempts to foster. Accepting the limitation of causality on the framework seems a lesser evil than the restrictions imposed by attempting to remove this general requirement. For specific models, the restrictions imposed by causality may be able to be relaxed on a case by case basis.

4.3.2.3 Stiffness and A-Stability

The time stepping method in turn can place some restrictions on the framework. Time stepping methods require a number of properties to be useful for a problem. For example, they should converge to the correct solution to the problem as $\Delta t \rightarrow 0$, and should be tolerant to the numerical errors

that are inevitably introduced by floating point calculations. Linear stability of the time stepping algorithms, and stiffness of the system are important things to consider when selecting a method.

Stiff systems are loosely defined as systems that require excessively small time steps for at least part of a simulation using some time stepping method (see [Iserles 2009](#), page 56). An example of how this can occur is if the solution to the system has a dominant component that evolves slowly, and a nondominant component that evolves much more quickly which does not contribute to the response of the system in an obvious way. For some algorithms, the time step size is governed by quick dynamics to insure that they do not introduce instability, whereas other algorithms can reject any disturbance from the quick dynamics to capture the slow dynamics accurately, despite using a time step set for the slow dynamics. The Runge-Kutta 4th order algorithm is an explicit method that is known to perform poorly on *stiff* systems (discussed below).

Every time stepping algorithm also has a certain linear stability domain. If all the eigenvalues (poles) of a linear system are located on the left half of the complex plane then the system itself is stable. It is desirable if a time stepping algorithm produces a stable solution for any system that has stable eigenvalues. This property is called *A-stability*. Unfortunately, no explicit Runge-Kutta method is A-stable. Instead, they have a stable domain that is a function of the time step size. This means that the time step size needs to be selected to ensure that all the poles of the system lie within this stable domain. (In cochlear models, we sometimes allow parts of the cochlea to be locally unstable. The remainder of the cochlea or nonlinearities present locally ensure that the system as a whole is stable.)

In this discussion, we have noted that explicit Runge-Kutta algorithms are not A-stable and perform poorly on stiff systems. There are other alternatives that do not have these drawbacks, and even offer other benefits such as preserving certain structure that might be present in the differential equations. The Runge-Kutta 4th order method has been widely used to date because it has been suitable for the models studied. When explicit methods are suitable, they are computationally cheaper than the alternatives. [Kim and Xin \(2005\)](#) and [Moleti *et al.* \(2009\)](#) address a cochlear model where explicit methods are not suitable, and both papers select implicit methods instead.

4.3.3 Strengths and Weaknesses

Having provided definitions of state space methods and the Runge-Kutta 4th order algorithm, and discussed some of the formal limitations on these methods, we now consider their strengths and weaknesses for a cochlear model.

4.3.3.1 Multi-Input-Multi-Output Systems

The methods extend naturally and easily to multi-input-multi-output systems. The process requires a sufficient number of state variables to describe the whole system. The effect of each input on the state, and the state variables on each other must be quantified. Finally, the effect of the inputs and the state on the outputs must be described. While the amount of work required to do this increases with the size of the system, the process is not fundamentally different from the single-input-single-output case. Given that the cochlea is a multi-input-multi-output system with a large number of components, this flexibility is valuable.

4.3.3.2 Nonlinear Systems

The methods handle nonlinear systems elegantly. The original definitions above consider the nonlinear case. In the linear case, the relationships between inputs, states, and outputs can be expressed in terms of matrices. For the nonlinear case, matrices are not sufficient, and nonlinear expressions relating the states to the inputs and outputs replace them. These methods allow a wide class of nonlinear equations to be expressed easily. The complication lies not in the representation, however, but in the time stepping algorithm. Stiffness of a dynamical system has been described above. Nonlinear equations can also be stiff which means that the explicit Runge-Kutta method would need to be replaced by an alternative time stepping method suitable for the equation in some cases.

4.3.3.3 A Body of Applicable Literature

State space methods have a body of literature and applicable techniques in the field of Control Theory. The theoretical problems that have been considered tend to be restricted to models that are much smaller and less extensively coupled than a cochlear model, but the significant difference is that the field of Control Theory aims to design controllers that modify the behaviour of the system, whereas a cochlear model aims to recreate what is present. Therefore, while a controller for a multi-input-multi-output system quickly becomes difficult to design, we are not immediately concerned with designing a controller, and hence the statement that state space methods extend easily to multi-input-multi-output methods remains true. A cochlear model is generally a simulation problem rather than a control problem, but having the tools of Control Theory available is nevertheless an advantage because they can be applied to parts of the model, and sometimes even to the whole system, to test its performance. (Elliott *et al.* 2007 used this argument as an advantage of state space methods.)

4.3.3.4 Time Delays

State space methods also have certain weaknesses. In particular, time delays are difficult for the theory in general. Time delays can be an important feature of cochlear models, because there is an

inevitable delay in the neural circuitry that modifies the cochlear behaviour in response to stimuli. [Ruggero and Rich \(1987\)](#) estimate that the delay of synaptic transmission in the inner hair cell alone is about 1 ms. The model presented in [Flax \(2008\)](#) relies on the delays expected in this feedback path. [Zweig \(1990\)](#) provides an analysis of a model with two separate time delays: a fast positive feedback loop; and a slower negative feedback loop. [Geisler \(1993\)](#) included an outer hair cell model with delay based on the physiological data available at the time. (A later paper by the same author indicated that more recent data did not show the same delay, prompting adjustments to that model, see [Geisler and Sang 1995](#).) This is another case where the fact that we are interested in simulating a system rather than controlling it becomes important. Time delays can be handled by buffering the values of signal for a certain number of time steps before using them, because we are only interested in the simulation. The problems associated with designing a feedback control system in the face of time delays are not of immediate concern in a cochlear model.

4.3.3.5 Existence and Uniqueness

The discussion of existence and uniqueness of the solution to a model should not be taken as a weakness of these tools *per se*. These are difficulties that are intrinsic to the model rather than the representation and implementation of the model. The discussion on causality is relevant here because there are alternatives that allow noncausal models to be implemented. We have argued, however, that restricting the framework to causal models in general is preferable to the restrictions that allowing noncausal models would impose.

4.3.3.6 Simulating Small Signals

Jont Allen pointed out that state space methods are not good at simulating small signals buried in large signals accurately (personal communication). His concern is that the distortion products that arise due to nonlinear processes in the cochlea are as much as 40 dB below the primary signal (see [Kim et al. 1980a](#)) and so care must be taken to simulate them accurately. This concern is noted here as a weakness of state space methods, and is taken up in the next section on alternatives to the state space approach.

4.3.4 Alternatives

4.3.4.1 Convolution

The final weakness of state space discussed above was that small signals buried in larger signals can be problematic. This is because the contribution of previous inputs is only recorded in the state of the system, which is updated iteratively. Convolution of the input with the impulse response of the

(linear, time-invariant) system is an alternative to solving the state space model using a technique like the Runge-Kutta 4th order algorithm. Convolution can be defined as follows:

$$y(t) = u(t) \star f(t) \quad (4.29)$$

$$= \int_0^t u(\tau) f(t - \tau) d\tau \quad (4.30)$$

$$\cong \Delta t \sum_{m=1}^n u(m\Delta t) f(n\Delta t - m\Delta t), \quad (4.31)$$

where $y(t)$ is the output, $u(t)$ the input, and $f(t)$ the impulse response of the system. The expression assumes $u(t) = f(t) = 0$ for $t < 0$ (Stremler 1990, page 114). That is, this representation assumes causality, but this is not required in general for a convolution. Δt is the time step such that $t = n\Delta t$, and m, n are integers. In the convolution, the contribution of each input is computed separately at each time step. This may improve the accuracy for small signals as discussed above, but it comes at some computational cost⁴.

The state space model for a linear single-input-single-output system with four states requires 204 floating point operations to compute a new state, and 9 more to compute the output, bringing the total to 214⁵. This number is independent of the time, but scales as n^2 with the number of states. Now consider solving the same system by convolution. The impulse response of a linear time-invariant system with two degrees of freedom (four states) is nonzero for all $t > 0$. The first time step requires just 2 floating point operations (both multiplications), but the N^{th} requires $N + 1$ multiplications, and $N - 1$ additions for $2N$ operations in total; see equation (4.31). This value increases with time, and to compute N time steps requires, on average N floating point operations. Assume that we want to simulate 100 ms, and let $\Delta t = 2.5 \mu\text{s}$. This means that on average 40 000 floating point operations per time step are required to calculate $y(t)$ – a significant increase in the computational requirements⁶. This requirement can be reduced by truncating the impulse response for $t > t_1$, where $f(t), t > t_1$ is considered to be suitably small. However, this has repercussions on the accuracy of the simulation as it introduces errors beyond those already introduced by the spatial discretisation. (Matthews used a similar technique to truncate the Green's function that he used for spatial integration in his PhD thesis. See his figure 2.3 for an interesting example of the system's behaviour for truncation at various points, Matthews 1980.) Techniques like parallel computing could be used to offset the

⁴However, experiments with simulating the response of a 2 state radial segment model to a sine wave using both state space and convolution suggested that for the same time step size the state space model provided a significantly lower L_2 error (measured against the analytical response). The error in the convolution case was dominated by integration error.

⁵Equations (4.20, 4.22, 4.24, 4.26) each require 36; (4.21, 4.23, 4.25) require 8 each; (4.27) requires 28; and (4.28) requires 8 floating point operations.

⁶A reviewer noted that discrete convolutions are often implemented using fast Fourier transforms, that have computational complexity of $O_1(n) = n \log_2(n)$, a significant improvement over the $O_2(n) = n^2$ of discrete convolution for n time steps. To minimise restrictions on the cochlear partition model, the time steps should be evaluated sequentially, providing the opportunity to compute new inputs to the model after each time step. Under these conditions there is no advantage to using fast Fourier transforms as the convolution equation is already $O_3(n) = n$. However, for more restricted partition models they may be viable.

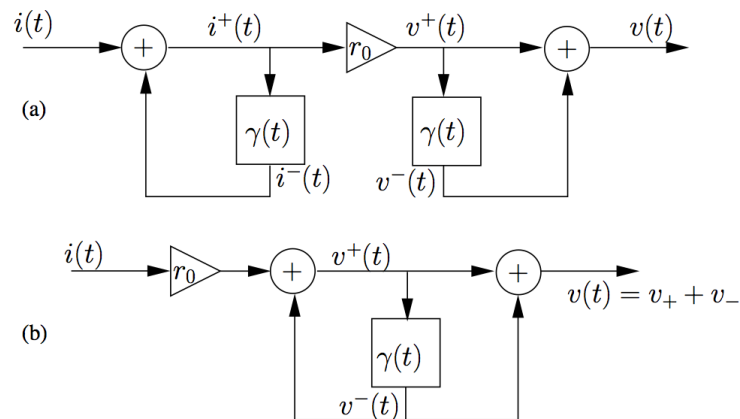


Figure 4.2: Two alternative methods of implementing the $V(s) = Z(s)I(s)$ in the time domain. r_0 is the surge resistance and is equal to z_0 . $\gamma(t)$ is the time domain “reflectance” kernel. $\gamma(t)$ has Laplace transform $\Gamma(s)$ defined by equation (4.32). $v^+(t)$, $i^+(t)$ and $v^-(t)$, $i^-(t)$ are the forward going and backward going current and voltage respectively. In a cochlear model, they could represent these components of the pressure and acceleration or the potential and velocity respectively (see section 4.2.4). By definition, $v(t) = v^+(t) + v^-(t)$ and $i(t) = i^+(t) - i^-(t)$. The two methods differ in that (a) keeps $i^+(t)$ and $i^-(t)$ as explicit signals, whereas (b) uses the fact that operators commute to eliminate these variables. (From Allen 2009, figure 1.)

additional expense of convolutions (see section 4.7).

Convolution can be extended to nonlinear systems (van Hemmen *et al.* 2000). Alternatively, a nonlinear system can sometimes be separated into linear blocks with simple nonlinear blocks. If this is possible, the nonlinear blocks can be implemented directly, and equation (4.31) can be used to implement the linear blocks.

4.3.4.2 Implementing Impedance in the Time Domain

Jont Allen has noted that the frequency domain concept of *impedance* is important in electrical engineering, partially because of its close link to power. He has proposed that models based on *reflectance* provide a time domain equivalent of impedance with similar properties (Allen 2009⁷). For an impedance $Z(s)$, parameters for its time domain implementation can be calculated using:

$$Z(s) = z_0 \frac{1 + \Gamma(s)}{1 - \Gamma(s)}, \quad (4.32)$$

where z_0 is a real number called the plane-wave characteristic impedance, and $\Gamma(s)$ is the plane-wave reflection coefficient, or reflectance. Figure 4.2 shows block diagrams of how these definitions can be used to implement an arbitrary impedance in the time domain. Panel (a) shows how the current and velocity, $i(t)$ and $v(t)$, may each be treated as the result of a forward going and a backward going

⁷This article is currently unpublished, but is available from its author at the website listed.

part (indicated by the superscripts $+$ and $-$ respectively). This factoring proved particularly useful in the middle ear model by Parent and Allen, where it allows the forward going signals and the reverse going signals to pass through separate time delays (Parent and Allen 2007). The middle ear has a significant time delay, but also allows otoacoustic emissions to travel from the cochlea back to the ear canal. If the reflectance representation was not used in this model, a high order model would be required to replace the time delay terms, which leads to its own complications.

A similar idea has also been used by Neely and Allen to assess the retrograde waves in a 1D cochlear model (Neely and Allen 2009). The pressure wave P in the fluid was explicitly separated into forward going P^+ and backwards going P^- components. In this case, they assumed that P^+ was not significantly affected by P^- to obtain a solution. They note that this assumption is equivalent to the WKB approximation. For the 1D cases considered in these two applications, the physical significance of these waves is clear. For 2D or 3D models it might also be possible to divide the fluid pressure into P^+ and P^- fields, but the meaning of these signals is not obvious. (It also appears that there are insufficient boundary conditions on these fields, but an assumption like that used in Neely and Allen 2009 may resolve this.) Nevertheless, the reflectance model may be used to simply replace the model of the cochlear partition (having taken nonlinearities into account) rather than attempting to obtain a complete solution in terms of P^+ and P^- .

A benefit of the principle of encapsulation that is at the heart of object orientated programming is that a framework for computational cochlear models need not select one approach to time integration above another. People who use and extend the framework are free to use any method that they prefer to implement an object, as long as the object accepts the same function calls, and provides the expected output for given inputs. The time integration methods discussed here will be required in the cochlear partition model. That model in turn interacts closely with the fluid model, as described in chapter 5. Therefore, the major constraint on the time integration methods selected is whether it can provide the specified interface for this interaction.

To date, only the state space model has been used for time integration in this framework, but the benefit of an object orientated framework is that when a class offering another variant is implemented, its functionality can be inherited by future classes. In an open-source project, this means that others can use the implementation, and even improve it, without needing to start from scratch.

4.4 The Cochlear Partition

Having discussed the methods available for implementing models of the cochlear partition in some detail, we return to the question of what characteristics the anatomy and physiology of the cochlea dictate that our framework must permit. Figure 4.3 shows a simplified cross-section of the organ of Corti complex that recalls the key features to be modelled. A recurring theme of this thesis in

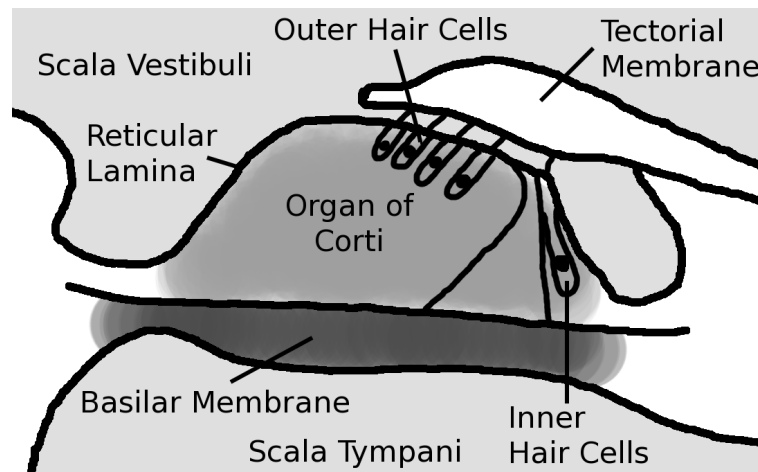


Figure 4.3: Diagram showing a simplified cross-section of the organ of Corti complex. Key features in the discussion are indicated. (Geometry corresponds to [Kessel and Kardon 1979](#), page 109.)

general, and this chapter in particular, is to ask the following question: “What is present in the cochlea, and how might we want to model it?”

Detailed Models and Lumped Parameter Models

Presented with the intricacy of even the simplified cross-section, and given that the finite element method has been selected for modelling the fluid domain, an obvious approach might be to use the finite element method to recreate this cross-section as accurately as possible. This approach has been taken by some researchers (for example, [Kolston and Ashmore 1996](#) and [Cai *et al.* 2005](#)), providing valuable insights into the micro-mechanics of the cochlea. Therefore, we need to consider the case that the cochlear partition model might be a finite element simulation in its own right. Models that aim to reproduce the cross-section in detail have certain weaknesses. For example, the model quickly becomes computationally intensive (this is related to the discussion on spatial dimension, section 4.2.6), and the model requires more physiological data to constrain the behaviour. Where the required data is not available, assumptions must be made. Lumped parameter models (sometimes called zero-dimensional models) are on the opposite side of the spectrum. They require a minimum of tuning parameters, many of which will be affected by multiple physiological features, and they are comparatively cheap to use in computation. State space models, discussed in the previous section, are a means by which lumped parameter models may be implemented.

Both of these types of models have a contribution to make to our understanding of the ear. Designing a lumped parameter model requires conscious decisions about the relative importance of features in the cochlea for its proper function, as well as their correct mathematical expression, and computational models that aim to recreate the physiology allow us to observe the behaviour of the cochlea in ways that are difficult to reproduce by experimental methods.

Ironically, both models can also obscure important details. For example, the active mechanisms in the cochlea have been included in lumped parameter models by allowing a damping term to be negative. The details of how this can occur need not be specified, which can simplify initial testing. However, if the benefit of the negative damping has been shown, we then require a more detailed model of its action. Satisfaction with the behaviour of the lumped model could slow down an investigation into its causes. Furthermore, lumped parameter models may also obscure the fact that important constraints which the cochlea works under have been ignored. For example, an active lumped parameter model need not include any of the passive dynamics. Active mechanisms work with, and to some extent also against, the passive behaviour, so neglecting the passive dynamics might affect estimates of the feasibility of a given active model based on the simulations.

Equally, when designing a detailed model two mechanisms that could contribute to a given behaviour might both be included (say compressibility and viscosity of the fluid, to draw on an example from the discussion about the fluid). When we look at the results we can see new behaviour. We think the behaviour is correct because of the detail in our model, and its good correspondence with the experimental results. However, in this example we are no closer to understanding whether the compressibility and viscosity of the fluid were actually important, and if so, how important.

The focus of this thesis is on lumped parameter models because, in the author's opinion, they offer a lot of scope for advancing our understanding of the ear with currently available computation power and physiological data. We anticipate that future lumped parameter models will include more details, and that detailed models will become increasingly important. Therefore, this framework should at least consider detailed models. The discussion of fluid-structure interaction in chapter 5 includes an investigation into a partitioned time integration scheme that is hopefully a step in this direction. Furthermore, it is argued below that the behaviour of passive lumped parameter models can serve as a template for the interaction between the cochlear partition and the surrounding fluid. Therefore, the changes required to extend support from lumped parameter models to detailed models might be surprisingly minor.

Location of Input and Output Boundaries

The scala vestibuli and scala tympani are sometimes modelled separately, but in the majority of models, a single representative scala is chosen (see section 4.2.7). If a single scala is considered then the pressure field on the boundary of the cochlear partition represents the pressure difference that acts on it. The reduction to a 2D model of the pressure field can be done by averaging values over, or by considering certain modes in, the omitted dimension (see [Viergever 1980](#), appendix 2c and [Ramamoorthy et al. 2007](#) respectively). The velocity on the scala tympani and scala vestibuli boundary is assumed to be the same for a single scala model. The choice of location for values on the scala tympani boundary is limited to the basilar membrane. In the scala vestibuli more options are available.

Some models only treat the basilar membrane, in which case the pressure and velocity would be theoretically located there. The reticular lamina and the upper surface of the tectorial membrane are also sometimes used as locations for the pressure and velocity on the scala vestibuli boundary. Obviously in a 1D or 2D model the geometry of these interfaces are simplified considerably. In the case of a single scala, these choices are largely implicit.

When two scalae are modelled, the choice is not as implicit because we do not require the velocity of the scala tympani boundary (the basilar membrane) and the scala vestibuli boundary (as selected) to be the same. If we model the organ of Corti as incompressible, and select the reticular lamina as our second boundary, then the velocities of both boundaries would be the same. Some cases when they would not be the same are as follows: if we model the organ of Corti as compressible; allow the outer hair cells to modify the behaviour of the reticular lamina in isolation via some mechanism; or allow the tectorial membrane to have its own dynamics, and select it as our boundary.

Radial Segment Sub-Objects

Why do we select the cochlear partition as a whole for a computational object rather than, say, selecting the tectorial membrane, organ of Corti and basilar membrane as separate objects? It is certainly consistent with the object orientated methodology to treat the cochlear partition as an aggregation of smaller parts, although the form of most cochlear models makes it more convenient to divide it along its length first (into radial segments). Sub-objects of these lengths may be introduced, but the lumped parameter models that are the focus of this thesis tend to incorporate these parts into a single dynamic model instead. If a separate object is required for the tectorial membrane in particular, the cochlear partition class is likely to be suitable to extend for this purpose, because the final specification of this class is very general.

We have discussed some of the effects of the anatomy of the cochlear partition, as well as some modelling considerations. In the remainder of this section we follow chapter 3 by dividing the discussion into passive dynamics, active mechanisms, coupling between radial segments, and innervation.

4.4.1 Passive Dynamics

4.4.1.1 A Simple Partition Model

Chapter 3 discussed the active and passive behaviour of the cochlea, and noted that passive behaviour is in many ways more a useful concept than an objectively measurable state of the cochlea. That discussion associated passive behaviour with the linear response that is observed when active mechanisms are not present, or are overwhelmed. The passive dynamics of the cochlear partition are dominated by its mechanical properties. Early studies on cochlear modelling justified models of the

type:

$$\Delta p_{BM}(x,t) = M(x) \frac{d}{dt} v_{BM}(x,t) + R(x) v_{BM}(x,t) + K(x) \int_{t_0}^t v_{BM}(x,t) dt, \quad (4.33)$$

where $\Delta p_{BM}(x,t)$ is the instantaneous pressure difference across the basilar membrane at a location x along the length of the basilar membrane⁸, $v_{BM}(x,t)$ was the velocity normal to fluid boundary at that point, and $M(x)$, $K(x)$, and $R(x)$ are the mass, stiffness, and damping respectively. Equation (4.33) can be derived from first principles by considering the forces acting on a radial slice of the cochlear partition. If this is done, it is clear that the physical parameters are in “per area” units, and that the breadth of the basilar membrane has dropped out. Furthermore, a sign convention has been assumed for the direction of positive $\Delta p_{BM}(x,t)$ and $v_{BM}(x,t)$. These are details that need to be taken into account when matching physiological data and implementing the model. It has been assumed that all the physical parameters depend on the x location. This dependence is dictated by the anatomy and cochlear frequency map, although in some types of models it is convenient to assume $M(x) = M$.

The certain structures have generally been assumed to contribute to each of these terms: $K(x)$, the basilar membrane; $M(x)$, the fluid in the organ of Corti; and $R(x)$, viscous effects in the organ of Corti. Early papers also included a longitudinal stiffness term, but physiological data (such as the direction of the stiffening fibres in the basilar membrane, and experiments deflecting the basilar membrane with a probe), as well as numerical results suggested that this term was not required (Allen and Sondhi 1979). In an interesting twist, a few recent models have included it again with some success. This thread will be picked up again section 4.4.3.

4.4.1.2 A Radial Segment Class

The model in equation (4.33) is continuous with respect to x . It is convenient to divide the x dimension into small discrete radial segments such that Δp_{BM} and the mechanical parameters are roughly constant for the length of the segment for numerical algorithms. This approach is supported by anatomical details of the cochlea. It has already been noted that longitudinal stiffness is neglected in many models, therefore these segments need not be directly coupled. (The case where they are coupled is discussed below.) Furthermore, a radial segment containing a single inner hair cell and its outer hair cells is a sensible limit on the smallest unit that might act in concert. Defining a radial segment class fits well with the object orientated methodology, because the behaviour of a single segment can be described and encapsulated. The cochlear partition object is then simply a set of these segments with some geometrical relationship.

Selecting the radial segments as class, we can note some of the features that it requires. This list will be extended as the discussion goes on, but so far it will require a state space model for the passive dynamics, together with an appropriate time stepping algorithm. The class will also need to interact

⁸At that time the basilar membrane, the key stiffness element in the cochlear partition, dominated the discussion, and hence the naming.

with the fluid object in up to two locations (the scala tympani and scala vestibuli). The coupling at these points will be qualitatively similar, therefore the class could contain two fluid interface objects that are responsible for this interaction. If the fluid interface class is designed well, it need not change as we consider active cochlear models, or even detailed (as opposed to lumped) models. This can facilitate the extensibility of the model.

4.4.1.3 A Template for Fluid-Structure Interaction

The second order model of the cochlear partition given in (4.33) provides a good match to physiological data despite its simplicity. In fact, it provides an excellent template for the fluid-structure interaction in general. Given the generality of the models that this thesis considers, providing adequate support for this claim is not trivial. Rather, the remaining chapters will each contribute in their own way to the question. We will require the mathematical formalism of section 5.1 to establish the result shown in section 7.2.2 – that the influence of fluid coupling between segments is determined by the partition mass. Section 5.2 will show that key characteristics are present in the fluid-structure interaction, whether the partition model is linear or nonlinear.

The final key element we require is to show that the Laplace fluid models we use extensively in this work are also representative of more complex fluid models. Viergever's dimensional analysis of the cochlear fluid, reviewed in section 4.2, gives us a formal method for accessing when the Laplace equation is an accurate fluid model. It shows that the assumption that the fluid is inviscid is particularly accurate for high frequencies, and over short time scales. Section 5.3 shows the importance of capturing the fluid-structure interaction over these time scales. The upper bound on the frequency range where compressibility is negligible is inversely proportional to a critical length. Section 7.2.2 shows that the important segment to segment coupling in the cochlea occurs between segments in the same neighbourhood. Therefore, in both of these regards the Laplace fluid model captures the key characteristics of the fluid. It is important to establish that equation (4.33) and Laplace fluid models can serve as a template for the fluid-structure interaction in the cochlea, because it shows that if the fluid interface is useful for models with the form of equation (4.33), then it can be generalised further to higher order models or detailed models.

4.4.1.4 A Parameters Class

The cochlear frequency map, and hence the values $M(x)$, $R(x)$, and $K(x)$, are dependent on species, whereas (4.33) is a reasonable first approximation for the passive behaviour of the cochlea across a wide range of species. It is therefore convenient to have an object class that is dedicated to providing these parameters. Specifying that the parameters are obtained from a separate class gives the user the greater freedom. For example, they can use derived classes that compute the values based on an equation, or on data read from a file without changing the radial segment class.

4.4.1.5 Data States and Data Model

During the simulation we might want to log the behaviour of some of structures or quantities. Many of the quantities of interest can be calculated using an output equation (4.19) and the state of the system (from here on referred to as the *dynamic state*). By considering a case using the simple passive model in equation (4.33) and a potential based fluid model (4.16), we see that some important values of interest cannot be derived from the state of the system alone. Assume we want to log the pressure difference across the partition in the model. If the partition displacement, velocity, and acceleration were states of the system, we would be able to compute the pressure difference using an equation like (4.33). However, as we noted, the relationship between pressure and potential is $p = \rho d\phi/dt$. Substituting this relationship into (4.33) and integrating both sides with respect to time yields:

$$\rho\Delta\phi_{BM}(x,t) = M(x)v_{BM}(x,t) + R(x) \int_{t_0}^t v_{BM}(x,t) dt + K(x) \int_{t_0}^t \int_{t_0}^t v_{BM}(x,t) dt dt, \quad (4.34)$$

which shows that the states required to model the dynamics of the potential model are $\int_{t_0}^t v_{BM}(x,t) dt$ and $\int_{t_0}^t \int_{t_0}^t v_{BM}(x,t) dt dt$. The variables required to compute the pressure difference are not part of this minimal (2 state) state space model. We could introduce the pressure as an explicit dynamic state, although it is not required, but this is not desirable because the ideal derivative that it entails cannot be implemented due to the causality requirement. To introduce it as part of the dynamic model would mean that unnecessary numerical errors are created. Furthermore, increasing the number of dynamic states in this manner is an inefficient use of the representation, as will be seen in section 6.3.3. Instead, it is preferable to allow each radial segment to have a second state space model to allow us to calculate any additional variables we want to log⁹. This simple example shows that it is efficient to include a separate state space model for additional variables that do not affect the dynamics (we will call this the *data state*). The distinction that we draw between dynamic states and data states is that only the dynamic states may influence the physical dynamics of the radial segment, and hence its interaction with the fluid and other segments, whereas both the dynamic states and the data states may influence the data outputs.

4.4.2 Active Mechanisms

Today, the activity in the cochlea has a well-known source: the motile outer hair cells. However, the exact mechanisms that these cells use to modify the cochlea from its passive state to its active state are still subjects open to debate and research. Regarding this thesis, the openness of this debate is essential to justify our interest in a general, extensible cochlear model. If a single model completely

⁹If we are just computing the derivative of a signal, it is simpler, and can be more accurate, to use a backward difference.

satisfied researchers in the field, there would be no need to compare and contrast available models. As it is, a wide diversity of mechanisms has been, and are still, considered.

The net effect of the active mechanisms should be to sharpen the cochlea's response to the input frequency, while compressing the dynamic range of the output. There should also be a range of other associated phenomena including distortion products, two-tone suppression, and otoacoustic emissions (see chapter 3). All of these features can be used to validate a cochlear model using a particular active mechanism.

4.4.2.1 The Local Feedback Pathway

Some models of interest stress active feedback paths from a given radial segment back to *itself* via the outer hair cells. Considering only local feedback is convenient from a modelling point of view: a radial segment only requires its own internal data to update its state and output. The values that might be used as input for the local feedback path can include any of the internal states (for example, the velocity or displacement of the basilar membrane or tectorial membrane) as well as the pressures at the fluid interface. The feedback path is likely to introduce additional states to the radial segment, either to capture the electrical dynamics of the hair cells, or the mechanical dynamics of the cells and their supporting structures in the radial segment. This augmented state space model of the cochlear partition dynamics will typically be nonlinear due to the nonlinearities present in the hair cells, but can also be solved using an appropriate time stepping algorithm. Finally, the feedback will be inserted back into the original dynamic model in some manner, because the outer hair cells ultimately modify the response in a manner that is measurable at the basilar membrane.

The active model will sometimes need to be linearised at each time step, but certain types of nonlinearities can be handled directly. Section 5.2 will discuss which types of nonlinearities fit into each group. Ultimately, the nonlinearities affect the way that the state gradient $\dot{\mathbf{X}}(t)$ and the segment outputs are computed. The details of how this can be achieved in an object orientated framework is discussed in section 6.3.

The partition model needs to receive input from, and return an output to, the fluid. In section 4.4.1, the argument that the same mechanisms can be used for input and output in both passive and active cochlear models was outlined. This means that the same fluid interface class can be used in both active and passive models. In fact, if active models are supported, then passive models are automatically supported as well, because the set of linear functions is a subset of all nonlinear functions.

4.4.2.2 Active Mechanisms Modify Passive Dynamics

When designing an active cochlear model, it is beneficial if the passive dynamics of the cochlea are well represented. There is a close link between the active and passive behaviour in the real cochlea.

The passive behaviour of a cochlear model emerges and dominates as the effect of active mechanisms decreases (either due to the size of the input, or because we explicitly reduce the active effects by some experimental procedure for example). Control Theory also shows that directly cancelling the dynamics of a system (pole-zero cancellation) is problematic in practice. Therefore, Control Theory suggests that it is unlikely that nature would use a mechanism that amounts to this sort of cancellation in the ear. The outer hair cells are significantly smaller than the structures that contribute to the parameters of the passive dynamics, so to directly oppose the dynamics of the passive system also seems to be an inefficient use of energy that comes at a cost to the animal. Therefore, if we neglect the passive dynamics in active models, we may miss important details of the interaction between active mechanisms and the passive structures. Furthermore, we might not be able to correctly assess whether a mechanism is plausible or not, for example, by estimating the work required by the active mechanisms in processing the input.

We have stressed the view that active mechanisms should build on the passive behaviour of the cochlea, but for a general framework imposing any restrictions in this regard is counterproductive. Therefore, in the implementation of active mechanisms in the cochlear modelling framework, we simply aim to allow general nonlinear equations to govern the local behaviour of a radial segment.

4.4.3 Physical Coupling between Segments

4.4.3.1 Arguments for Including Coupling between Segments

If a radial segment models a section of the cochlear partition containing a few inner hair cells (with a lumped parameter approach the same model can represent a partition of varying length) then the assumption that the feedback paths are mostly internal is justified. As the length of the radial segment is decreased, the assumption becomes more tenuous. [Fernández \(1952\)](#) provides values for the cross-sectional area of the organ of Corti (at least 0.012 mm^2) and the width of the basilar membrane (0.2 mm) in the guinea pig. Therefore, the height of the organ of Corti alone is of the order of $60 \mu\text{m}$. The width of hair cells is of the order of $10 \mu\text{m}$ (see [Sato et al. 1999](#) for cat data, in humans a similar estimate can be obtained from the length of the basilar membrane, 35 mm , and the number of inner hair cells, 3000). If we are modelling a radial segment that is more than $60 \mu\text{m}$ high, but just $10 \mu\text{m}$ wide (corresponding to a single inner hair cell width) then the assumption that radial segments do not interact no longer seems reasonable. Given the surface area exposed to the fluid ($\approx 2 \times 10 \mu\text{m} \times 200 \mu\text{m}$) and the surface area exposed to other segments ($\approx 2 \times 60 \mu\text{m} \times 200 \mu\text{m}$), some relevant mechanical forces at these interfaces should be expected at this level of refinement, particularly given that estimates for damping in the partition ($R(x)$ above) are often attributed to viscous forces in the organ of Corti (for example, see [Mammano and Nobili 1993](#)).

Aside from these simple geometric arguments, the presence of longitudinal stiffness in the basilar membrane also implies that there is coupling between the radial segments. Early computational




Figure 4.4: Diagram from Geisler and Sang (1995), figure 1B, showing the arrangement of outer hair cells in a radial segment. Notice the basal slant of the hair cell itself, and the phalangeal process of the Deiters cell extending in the apical direction. Geisler and Sang (1995) assumed the relatively large value $\phi = \pi/4$ in tests. A scanning electron micrograph of the mole rat cochlea shown in Slepecky (1996), figure 2.8, suggests $\phi = \pi/10$ would be appropriate for the mole rat.

simulations suggested that longitudinal stiffness was not required, and that including it in models would have a negative impact on the quality of the fit to physiological data. More recent work suggests that it might be important. For example, Ramamoorthy *et al.* (2010) found that when treating the basilar membrane as a plate, including a longitudinal Young's modulus that is 10% of the radial Young's modulus improved the qualitative match between the frequency dispersion in physiological data from the apex of the cochlea and their model. Meaud and Grosh (2010) found that including viscoelastic coupling in the basilar membrane and tectorial membrane could improve the stability of cochlear models, and their match to experimental data. Therefore, while many models ignore longitudinal stiffness, there has recently been renewed interest in it.

Finally, the anatomy of the organ of Corti suggests that forces between rows of hair cells are deliberately encouraged and controlled. Geisler and Sang (1995), in an early report to include the effect, presented a clear schematic showing that the outer hair cells have a basal slant from their supporting cell, the Deiters cell. Their diagram is reproduced in figure 4.4. The Deiters cells also have a phalangeal process that joins the reticular lamina at an apical point. Therefore, the active forces developed by the outer hair cells affects at least two radial segments. This anatomical structure would also have an effect on the passive mechanics of the radial segments.

4.4.3.2 Implications of Coupling between Segments

The fact that there are mechanical forces between radial segments in the biological cochlea has an impact on the structure of the radial segment class. In the object orientated paradigm, each object encapsulates and owns certain data that describes itself (notably its state in the case of a radial segment). To implement forces between segments, each segment needs access to the dynamic state of its neighbours when updating its own dynamic state, so some mechanism needs provide the required data.

Aside from its dynamic state, each radial segment is also affected by the pressure (or potential) in the fluid channels above and below the segment. Can these values have a direct effect on the dynamics of the neighbouring segments? If they only have an indirect effect, via the dynamic state of the system, then it is sufficient to allow neighbours to access the state only. When determining the dynamic model of a radial segment, it is important to consider the forces that act on that segment. This is essentially how equation (4.33) is derived. If the coupling between the segments contains only stiffness and damping terms, then the force acting on the segments is determined by their relative displacement and velocity. Equation (4.33) shows that these will form part of the state. Even after transformation to a potential based fluid model, their contribution continues to be determined by the state (see equation (4.34) to confirm this). One way for the fluid input of one segment to have a direct effect on a neighbour would be if the force that it exerts is shared between the two segments. Given the slant in outer hair cells shown in figure 4.4, this kind of force sharing seems plausible. Alternatively, if the outer hair cells are sensitive to pressure, the slant of the hair cell will guarantee that at least two segments receive its direct input.

The literature provides some examples of the two basic types of coupling between segments: the model in [Geisler and Sang \(1995\)](#) only requires access to the state of neighbouring segments; whereas a more recent work allows fluid forces to act directly on neighbouring segments ([Kim and Xin 2005](#)). In fact, the direct force seen in [Kim and Xin \(2005\)](#) arises because outer hair cell forces, which depend on *state* of the segment, are converted to an equivalent *pressure*, and added to the input pressure of neighbouring segments. This suggest that the model does not have any *inherent* direct fluid pressure effect, and could be adapted to avoid this input, but it is nevertheless useful to be able to simulate the model as it stands. Therefore, a consideration of both the process of cochlear modelling in general, and some specific models, illustrates that it is important to provide access to the fluid inputs to coupled segments, as well as the state. Of the two, the access to the states will be of more general use and importance.

So far we have considered forces between direct neighbours. However, the physical coupling between radial segments may not be restricted to neighbours only. For example, the phalangeal process may extend a couple of segments apically. In generally, a similar mechanism to the one used to allow interaction between neighbouring segments can be used to allow interaction between arbitrary

pairs of segments. Assuming that these interactions will only occur between segments within some specified distance along the cochlear partition allows for optimisation in the implementation, in particular if the model is parallelised using message passing (see section 4.7). Based on the previous discussion, the assumption that coupling only occurs between segments in the same region is true to the physiology of the cochlea.

Therefore, providing a mechanism to allow interaction between radial segments requires that the segment class provides read access to its dynamic state and current fluid inputs. It also requires a mechanism for a specific segment to “link” to another segment (normally in the neighbourhood, but not necessarily) in order to read these values. The segment reading the data would incorporate the data into its state update equations, as required by the model.

4.4.4 Innervation

Sections 3.1.3 and 4.4.4 described the innervation of the cochlea. There are both afferent nerves that carry information to the brain and efferent nerves that return information from the brain. In addition to providing input to the brain, these nerves form multiple feedback paths that include the cochlea and the brain. Therefore, cochlear models may relate to brainstem models in two ways: they may simply provide input to the brainstem models for further processing; or they may also rely on processing pathways through brainstem models. Using a separate brainstem model as a subsequent processing stage or as an additional feedback block requires some conventions on the interface that each provides. To date however, there are still sufficient open questions surrounding each of these blocks in isolation that efforts to interface separate models have been limited. Instead, a number of cochlear models have assigned roles to brainstem feedback, but implemented this feedback within the cochlear model itself.

From the perspective of the radial segment class, innervation requires that the class provides some data to, and accepts some input from, external objects. We have already observed that a radial segment requires data outputs based on internal data states. This same mechanism may be used to provide afferent outputs, because the state space model can represent the ionic concentration in the inner hair cell just as easily as it can represent the pressure difference across the cochlear partition (to recall our previous example). (Liu and Neely (2010) and Ramamoorthy *et al.* (2007) are examples of models that assign states to ionic concentrations in hair cells.) The input from the external objects would need to be able to affect the dynamic state of the segment in order to close the mechanical feedback loop, although it could also have an effect on the data states. The potentially nonlinear state space models assigned for the dynamic states and data states are flexible enough to allow this, provided that these models make provision for external inputs.

Once again we need to check causality of the feedback path. In this case, the feedback need not be to the same area as the signal arose, so if the feedback path is not strictly causal it can permit instan-

taneous coupling between remote locations on the cochlear partition. Based on the physiology of the cochlea, the neural feedback paths must be strictly causal, because there is a synaptic transmission delay and nerve spikes propagate down a nerve fibre at a finite speed (Ruggero and Rich 1987). This allows us to restrict the extent of coupling that is allowed via the neural pathways to simplify the implementation of the framework. If causal, as opposed to strictly causal, feedback is required for specific models, note that the mechanisms which allow physical coupling between segments can be overloaded with the neural model, although the benefits (in particular for parallelisation of the model) of restricting the range of those mechanisms will be lost.

So far the discussion has focused on the feedback path via the brainstem. We now turn to consider the path for afferent data in more detail, and ask at what point along the transmission path we define the end of the cochlear model, and the start of the next model. Each radial segment has (in the limit) a single inner hair cell. The behaviour of an inner hair cell is governed by the concentration of certain ionic species, the deflection of its stereocilia, and neural feedback. Overall, it requires relatively few states to describe its behaviour. These states could influence, and are influenced by, other states in the dynamic model directly. Therefore, inner hair cells can reasonably be included in the cochlear model. The afferent nerves on the other hand are much more numerous than inner hair cells (based on 30000 afferent fibres for 3000 inner hair cells in man, there are about 10 times as many afferent fibres). Each of these fibres requires relatively little input from the inner hair cell (just the ionic concentrations) and can be represented by a state space model in its own right. Therefore, a logical place to draw the line between the cochlear model and the next stage of neural processing is between the inner hair cells and the afferent nerves.

4.4.5 Cochlear Partition Compilation

So far we have said relatively little about the cochlear partition class itself. Is it simply an array or list of radial segments? The radial segments that form the basis of this class each need to obtain input from the fluid object and return outputs. There is some computational overhead required to locate the segment with which each point in the fluid must interface. This becomes an even bigger factor when the diversity of allowable fluid models is considered. The cochlear partition class is a convenient place to put the code required to create segments, and pair them with specific fluid locations. In particular, it can initialise the parameters object, and pass it to each segment as needed. It also collects and distributes external inputs, as well as passing communications between radial segments. This removes the burden of these operations from the radial segment objects. Therefore, in this framework, the cochlear partition class is assigned a largely data management role, with the physiological operations modelled by its aggregation of radial segment objects.

4.5 Input Boundaries

The fluid plays the role of transmitting signals between remote parts of the cochlea. In addition to describing its behaviour within the fluid region, we also need to constrain its behaviour at its boundary to obtain a well-posed model. Some parts of the boundary simply apply a load to the fluid, but others allow signals to enter the cochlea from the outside environment. The cochlear partition boundary has been discussed in detail in section 4.4.

At the risk of stating the obvious, the function of the cochlea is to transduce signals from the outside environment into a form that the brain can process further (neural impulses). It is important to come back to this simple statement at this stage, because we have noted that an active nonlinear cochlea can generate harmonics in response to pure tones, and distortion tones in response to pairs of tones. Spontaneous otoacoustic emissions are evidence of tones generated in the cochlea that are not even the response of a nonlinear system to some external input. All of these generated tones can be observed in the ear canal. They are partially reflected at the stapes, and the psychoacoustic impression created by these sounds are generally indistinguishable from tones not generated by the cochlea. These tones may be viewed as “inputs” in a sense. Indeed, this approach was taken by Matthew Flax in his PhD work (Flax 2008). Therefore, we return to our statement above to remind ourselves that the brain is not primarily concerned about the state of the ear, but about the state of its environment¹⁰. A comprehensive physiologically accurate model should make provision for the input pathways to be modelled accurately.

Signals enter the cochlea via two major pathways: air conduction using the outer ear, middle ear, and stapes; and bone conduction. We consider each of these in turn. During this section, we consider what features the classes that model these boundaries would have to include. It goes without saying that input boundaries need to provide at least one function that allows the external input to the cochlea to be set.

4.5.1 The Outer Ear, Middle Ear, and Stapes

Air conduction is the main input pathway to the cochlea. The outer ear and middle ear efficiently couple inputs into the cochlear fluid. In a linear cochlear model, the stapes pressure or acceleration can safely be taken as the input to the model, because the model response contains the same frequencies as the input. When considering a nonlinear model, we must be aware that wherever we impose an input, we are forcing any off frequency components generated in the model to cancel at that point, and therefore any generated components cannot be assessed at that point. If our interest is in the behaviour of the cochlea itself, then we must ensure that the point selected for our input does

¹⁰In defining the normal input path to the cochlea strictly, we do not intend to criticise the approach just mentioned, which is a useful concept for the purposes of that work.

not impose constraints on the active mechanisms. In particular, enforcing a pressure or velocity at the stapes will cause the cancellation effect mentioned above.

Effectively, our model requires a few degrees of freedom between the point we impose our input and the point we observe the results to ensure that the observations are not constrained by the input. These issues are discussed further in [Kim *et al.* \(1980b\)](#). A practical outcome of this discussion is that our model of the air conduction pathway must contain an extra “block” beyond the point we are interested in reading the results: if we are interested in the behaviour of the stapes, we require a middle ear model; if the behaviour of the tympanic membrane is of interest, a model of the ear canal is required. In his PhD, John Matthews even included a earphone model to begin the conduction path ([Matthews 1980](#)).

Chapter 3 provided an overview of the types of models used in this pathway. With the exception of the wave model introduced by [Parent and Allen \(2007\)](#), all of these models may be represented and simulated by state space models. Therefore, a class handling this boundary would also require a state space model in general. The wave model requires the reflectance ideas discussed in section 4.3.4.

We now consider the fluid interface that this class requires. A simple model that assigns two degrees of freedom to the middle ear, such as the one used in [Neely and Kim \(1986\)](#), is equivalent to equation (4.33) with the pressure difference redefined to account for the pressure at the ear canal, the transformer ratio of the middle ear, and the back pressure on the stapes due to the cochlear fluid. The $M(x)$, $R(x)$, and $K(x)$ in that equation would be lumped parameters, without x dependence, based on a fit to the middle ear transfer characteristics. The fluid interface required for this class is similar to that required for the cochlear partition, and the argument that the cochlear partition interface should extend to higher order models applies equally well to the fluid interface of this class: if this simple middle ear model can be implemented, more advanced models can use the same implementation.

A significant difference between the fluid boundary at the stapes and at the cochlear partition is that the stapes moves as a unit, whereas the cochlear partition does not. When considering the cochlear partition, we divided the structure into individual segments that each move as a unit. For the stapes this is not required, and a single lumped model can describe it, but care needs to be taken when calculating the forces acting on it. In this case, the pressure of cochlear fluid acts across multiple finite element cells, and multiple cells require input from the model. The class requires an integration rule to compute the total fluid force acting on the stapes. A model might also include edge effects in the flow pattern near the stapes so a “rule” that allows the input to the fluid to have a spatial dependence is also required.

From an object orientated point of view, these requirements are actually the requirements on a stapes model. We could view the middle ear as an aggregation of other models, connected in series along the signal path. If this is done, care needs to be taken to check the causality of models along the path, and to ensure that the coupling between these blocks satisfies the causality requirements imposed.

4.5.2 Bone Conduction

The second input path for signals to reach the cochlea is via bone conduction. When the stem of a struck tuning fork is placed against the forehead, it produces a more distinct sound than the same tuning fork held slightly away from the forehead. In this case, the sound need not follow a single pathway. [Stenfelt *et al.* \(2004a\)](#) lists the following pathways:

- Sound radiated into the external ear canal and middle ear cavity is transmitted to the cochlea via the tympanic membrane and the middle ear ossicles;
- Inertial effects of the middle ear ossicles;
- Inertial effects of the fluid in the cochlea; and
- Compression and expansion of the petrous bone encapsulating the cochlea.

Capturing all of these various effects in a cochlear model requires modifications to various objects in the cochlear model. The first two concern the model of the air conduction pathway discussed above. The third requires a body force to be applied to the fluid model. The final pathway can be implemented by imposing a nonzero normal velocity at the bone wall boundaries. The models described so far are flexible enough to incorporate these effects.

4.6 Other Boundaries

Besides the boundaries that can serve as inputs to the cochlea, there are others that need to be represented in a physiologically accurate and well-posed computational model. These include the round window and/or helicotrema, as well as the petrous bone. The bone surrounding the cochlea is generally modelled as infinitely stiff, so a zero normal flow boundary condition is normally imposed when bone conduction is not studied. The oval window at the stapes and the round window are the two most significant areas where the outer boundary can deform allowing the geometry of the fluid region to change. However, there is also evidence that other mechanisms come into play, in particular for pathological conditions. These other mechanisms are collectively called the “third” window, and will also be discussed below.

4.6.1 The Round Window

The round window is a gap in the bone encasing the cochlea, covered by a flexible membrane. The role generally ascribed to the round window is a pressure reference, and release, for the cochlear fluid. It allows the nearly incompressible cochlear fluid to displace in response to stapes motion.

However, it has significant stiffness. Lynch III *et al.* (1982) quote the compliance for the round window as $10 \times 10^{-9} \text{ cm}^5 \cdot \text{dyn}^{-1}$. For comparison, they quote $0.37 \times 10^{-9} \text{ cm}^5 \cdot \text{dyn}^{-1}$ for the stapes annular ligament. While the stiffness of the round window is ≈ 30 times smaller than that of the annular ligament, they determined that it is the second largest stiffness contributing to the input impedance of the cochlea.

In cochlear models with just a single scala, the round window does not appear. It is implicitly assumed that the flow pattern around the round window is symmetrical to the oval window, but with the direction of flow reversed. The stiffness of the round window is either neglected, or incorporated into the stiffness component of the stapes or middle ear model.

For a cochlear model with two separate scalae, an object to represent the round window is required. The modelling requirements for the round window are in many respects similar to the stapes boundary, and the same dynamic state space model structure is adequate to represent the stiffness dominated dynamics. If the pressure in the middle ear cavity is not constant, the round window may even be an additional input path. There are two significant differences. Firstly, the round window only moves as a unit at sufficiently low frequencies (below 1.5 kHz, Stenfelt *et al.* 2004b), otherwise it can support more complex behaviour. (If it is sufficient, to a first approximation, to assume that it moves as a unit, then the input boundary object is adequate. Failing this, the round window could be treated by a model consisting of aggregated parts, similar to the cochlear partition model.) Secondly, the choice of free and prescribed variables at the boundary are generally opposite to the stapes. At the stapes it is common to prescribe the acceleration (or velocity) into the cochlea, and allow the pressure (or potential) to serve as feedback to the stapes model. This is also possible at the round window, but it is more natural to treat the pressure (potential) as the prescribed, and acceleration (velocity) as the feedback variable. This second approach leads to a causal transfer function of the form: $p_{ow}(t) - p_{mec}(t) = (K_{ow}/s^2) a_{ow}(t)$, where p_{ow} is the fluid pressure on the oval window, p_{mec} is the pressure in the middle ear cavity, s is the Laplace variable, K_{ow} is the stiffness of the oval window, and a_{ow} is its outward acceleration. The inverse of this transfer function is not causal in isolation, requiring nondominant poles for a causal representation. (In fact, this is because the round window also has mass and damping properties that would provide the required poles, but they have been ignored.) Swapping the prescribed and free variables requires changes to the procedures to integrate the input to the model and allow spatial dependence for the output. This variation of the fluid interface can be incorporated into the input boundary class, making an “other boundary” class that is suitable for either case. The case where the round window does not move as a unit requires an object similar to the cochlear partition object.

4.6.2 The Helicotrema

The helicotrema is an important boundary condition for cochlear models that use a single scala. However, as a boundary condition, it is introduced by the transformation from a two scalae model to a single scalae model. In fact, it is essentially a pipe connecting the scala tympani to the scala vestibuli, allowing these fluid chambers to establish an ionic equilibrium, and to transferring fluid borne waves between the two scalae. In the cat at least, the helicotrema is not the dominant constriction in the fluid path from stapes to round window: the scala tympani has a similar cross-sectional area in the apical region (Lynch III *et al.* 1982).

The significance of this fact is that many cochlear models include damping at the helicotrema boundary when reducing to a single scala. (Neely and Kim 1986 is an early example.) Dallos (1970) presented a study of the low frequency behaviour of the helicotrema. They gave an expression for the theoretical impedance of the helicotrema (modelled as a circular pipe of given radius and length). The expression consists of a mass term and a damping term, where the damping term is proportional to the fluid viscosity. In the paper, Dallos stresses the importance of the *mass* of the fluid in the helicotrema at the expense of the *damping*¹¹. Therefore, the justification for including damping at the helicotrema does not seem to have support from the literature. It seems inconsistent to neglect viscosity in the scalae fluid yet include it in helicotrema fluid, given the dimensions of these fluid spaces. Finally, Puria and Allen (1991) found that if other physiological details were incorporated into the cochlear model (notably: scalae tapering; a realistic cochlear frequency map; and viscosity of the fluid in the scalae) then the details of the helicotrema boundary condition had little effect on the (input impedance of the) model. Specifically, reflections at this boundary were minimal.

Apical reflections are the reason that researchers include damping in helicotrema models. As we already noted, they were observed in the inviscid cochlear models by Matthews (1980) pages 70–75 and Neely (1978) page 56. If the impedance of the helicotrema is well matched to the impedance of the cochlea's apical end, then these reflections are minimised. Incorporating a damping term allows this to be done. While this damping term may not have direct support from the literature, strong apical reflections in a cochlear model are nonphysical behaviour – implying that the damping is a useful tool. In this author's opinion, it is important to note that by incorporating damping we are effectively accounting for approximations made at other points in the model, specifically concerning fluid viscosity and scalae tapering. We also have alternatives to including damping: scalae tapering has been considered in some models; viscosity can be included in finite element fluid models (at some computational expense, see appendix E); and the cochlear frequency map can also be manipulated to avoid apical reflections (Puria and Allen 1991).

Having discussed the options for a helicotrema model in some detail, the requirements that they place on the framework can be dispensed with relatively quickly. In a two scalae model, the helicotrema

¹¹In this author's opinion. Given that this differs with common modelling practice, the reader is encouraged to form their own opinion from the report.

can be included in the fluid region, and no separate object is required. (Some means to ensure that apical reflections are not problematic might be required.) Alternatively, an explicit helicotrema object may be included if desired. In some tests of the framework the cochlear partition class was used to provide a helicotrema object, because it too has a fluid interface with both the scala vestibuli and scala tympani. For a single scala model, the “other boundary” class is suitable, because the various conditions that might need to be imposed at the helicotrema are similar to those at the round window, which have already been discussed.

4.6.3 The “Third” Window

The fluids in the cochlea are maintained by the vestibular aqueduct (entering the vestibule) and the cochlear aqueduct (entering near the round window) respectively. These, and other structures, can serve as additional boundary conditions on the fluid in the cochlea, but they are generally neglected because they are less significant than the boundaries already considered. In pathological cochlear conditions, their contribution can become significant. In experimental work, Tonndorf and Tabor tested the effect of closing the oval window and the round window ([Tonndorf and Tabor 1962](#)). With the cochlear aqueduct unobstructed, they found similar shifts in the response amplitude for a sealed oval window and those with both windows sealed for bone conduction stimulation of the cochlea. When they blocked the cochlear aqueduct, closing the round window as well as the oval window caused an additional shift in the response amplitude compared to the case where only the oval window was blocked. This shows that the cochlear aqueduct can influence cochlear behaviour when the other windows are immobilised. In less extreme experiments, the effect of the third window can be assessed by measuring the difference between the fluid flow at the two windows. [Stenfelt *et al.* \(2004a\)](#) did this type of experiment, and found that the values could differ by as much as 15 dB for bone conduction stimulation. For air conduction stimulation, the values differed by less than 3 dB.

“Third windows” may also be directly caused by pathologies, rather than simply being brought into play by the pathology. [Merchant and Rosowski \(2008\)](#) classify the “third windows” resulting from various medical conditions as discrete or diffuse, and describe the relevant conditions that they arise.

Cochlear models that study phenomena related to the third window would need some mechanism for applying these types of boundary conditions. The discrete third windows may be incorporated into a cochlear model using an “other boundary” object. Diffuse third windows would require an object similar to the cochlear partition object. If there are compressible elements in the cochlear partition, these could also have effects similar to third windows, for example it would also allow the fluid volume displacement at the stapes and round window to differ.

4.7 Parallelisation

What is Parallelisation?

Parallelisation of computer programs is becoming an increasingly important way to reduce the run-time of the programs. This is because it has become inefficient to increase the frequency of computer processors in an attempt to obtain more floating point operations per second (flops) from the processor. Instead, modern computer processors run at more moderate frequencies, but provide multiple *cores*, each capable of executing programs independently (Olukotun and Hammond 2005). Aside from multiple cores on a single chip, multiple chips can be connected to the same bus on a computer mother board. Separate computers may even be connected together via a network interface to pass the data and instructions for a single program between them. These strategies allow the overall flops available for a computation to continue to increase, but they require a paradigm shift in the programming approach.

There are two major approaches to parallelisation: MPI (message passing interface); or SMP (shared memory processor), see Gropp *et al.* (1999). Shared memory processors provide every core with access to all the memory available. The details of how this is achieved are hidden from the user, and this type of machine is easier to program because the data structures do not need to be parallelised, but there is overhead associated with this type of machine, and these machines tend to be relatively small. Message passing can be used on shared memory machines, but comes into its own for distributed memory machines. Distributed memory machines require simpler hardware, but only provide access to the locally available memory. Instructions and data are passed between the cores via the message passing interface. While these machines are harder to program in general, the simpler hardware means that the largest clusters only provide this type of environment.

Currently, writing software to fully utilise these machines is challenging (Sutter and Larus 2005). We need to find opportunities for parallelism in our problem. A governing principle in parallelising a computer program is to exploit parallelism at the highest level first: when parallelising a nested loop, for example, it is more efficient to parallelise the outer loop than the inner loop. Parallelisation is essentially an optimisation step in that it must not change the result of the program. Like other optimisation steps, it is best to first direct our energy at the portions of code that contribute most towards overall simulation run-time.

Opportunities for Parallelism in Cochlear Models

In view of these principles, cochlear models offer many opportunities for parallelisation. It is easy to overlook the best location for parallelisation in a cochlear model: performing individual simulation from a batch of simulations in parallel. Physiological experiments often require a number of separate tests on each subject (for example testing the response to a set of different input frequencies). These

must also be performed on each subject in the sample. If the same set of tests was performed on a time domain model, they could all be run in parallel on separate machines, because they do not require any input from each other during the simulation. Therefore at the level of a batch of simulations, cochlear modelling is *embarrassingly parallel* – the term for problems that divide into parts having no interdependence.

The second natural place to look for opportunities to parallelise the cochlea is in the spatial and temporal grid used to simulate the cochlea. We discretise the spatial dimensions using the finite element method, and the time dimension using the time stepping algorithms covered in section 4.3. Consider the parallelisation of the spatial dimension: If we assume that the maximum number of radial segments we might want to simulate is 3000 (based on the number of inner hair cells in man), and that the ratio of the height of the cochlea to its length is about 1:30 (the human basilar membrane is about 35 mm long and the cochlea is about 1 mm high on average in man), then the number of spatial degrees of freedom is about 3000×100 , if we assume that the finite element mesh uses square cells and a single scala in a 2D model.

In a simulation, the finite element method requires two main steps: assembling the matrices for the problem; and solving these matrices to obtain a solution. The assembly procedure parallelises very efficiently in general because the process involves separate cells. deal.II provides various options for parallelising assembly: thread-based parallelism using pthreads; task-based parallelism using Intel Threaded Building Blocks; and message passing threads are all provided. Two main libraries provide solvers for a message passing parallelisation model in deal.II: Trilinos, maintained by Sandia National Laboratories ([Heroux et al. 2003](#)); and PETSc, maintained by Argonne National Laboratory ([Balay et al. 2011](#)). Matrix multiplication operations are also automatically parallelised in deal.II, so on shared memory machines the solving stage is automatically parallelised to some degree (only a limited number of parallel preconditioners are currently provided).

The cochlear partition is also spatially discretised, which suggests that the job of simulating the radial segments could be divided over a number of cores. If there is no physical coupling between the radial segments, then this process is *embarrassingly parallel*, but care must be taken if physical coupling is present. In the case of a shared memory machine, all that is required is to ensure that a radial segment does not overwrite a value that a linked segment still requires (since ownership of the shared values is clear). In a message passing model, the variables required by one segment may be in the local memory of another core. Therefore, we need to arrange to pass these values between the two cores. If there is some restriction on the spatial range of the physical coupling forces, then this can simplify the process of ensuring that all values that might be required are made available. The physical justification of this restriction was discussed in section 4.4.3. Parallelising the cochlear partition is a useful way to manage the increased computation burden as the radial segment model becomes more detailed. In particular, in section 4.3.4 it was noted that convolution might provide a useful alternative to state space models with time stepping algorithms, but it can be considerably

more computationally expensive. Computing these convolutions in parallel is a viable way to reduce the impact on the run-time of the simulation.

We can also consider dividing the problem into sets of time steps for each core to perform. There would be two main reasons to consider this: if there is less dependence between the jobs performed by each core when divided by time step; or if there are significantly more time steps than spatial degrees of freedom. Both of these factors imply that less data needs to be passed between the various cores. Cochlear simulations generally require time step sizes of about $2.5 \mu\text{s}$. Therefore, in a simulation of $2.5 \mu\text{s} \times 300 \times 1000 = 0.75 \text{ s}$ the number of time steps required is the same as the number of spatial degrees of freedom. If the length of the simulation is much larger or the time step size is much smaller than these assumed values, then parallelising the time axis first would seem to be the better option. Unfortunately, to parallelise the time axis assumes that we have a good understanding of the temporal relationships that exist in the model. In fact, the specification allows the boundary conditions (in particular for the cochlear partition) to be nonlinear, and even allows for undefined external feedback loops. The flexibility of the feedback blocks that we have motivated makes it impossible to provide a general parallelisation of the time axis. Specific models could sacrifice the generality of the nonlinearities and feedback to allow the time axis to be parallelised. Overall, however, parallelisation of the time domain is not possible in a general cochlear modelling framework of this type.

Parallelisation and 2D Laplacian Fluid Models

In the face of the global shift from sequential to parallel programming, it came as a big surprise to find that the type of cochlear models tested in this thesis benefit little from parallelisation. Appendix B shows results of tests performed on the model studied in chapter 7. Solving the fluid model required the majority of the software run-time. In the tests, sequential solvers outperformed parallel solvers even when the odds were stacked in favour of the parallel solver. This is because in a 2D model the largest spatial problem is too small to fully reap the benefit of the parallelisation. Rather than investing energy into parallelisation, the code used to date is essentially sequential, and multiple tests have been run in parallel to make use of the *embarrassingly parallel* nature of separate experiments. This is not to say that parallelisation of cochlear models is not important: it will be essential for the detailed models that we have suggested will inevitably be required (hence the extended discussion in this chapter). Rather, the results and discussion in appendix B show why parallelisation of the model was investigated, but then neglected for *this* work, in spite of the rapid trend towards parallel programming world wide.

4.8 Summary

This section has introduced the major classes required to simulate the cochlea in an object orientated framework: a fluid model; a cochlear partition class (consisting of radial segment objects); and an “other boundary” class. The discussion of the possible fluid simulation requirements led to the conclusion that the finite element method is more than adequate to simulate the fluid (see section 4.2.9). In this section, we provide a summary of the requirements for the remaining classes.

Figure 4.5 shows the cochlear partition inheritance diagram. It is modelled by a class named `OrganOfCortiComplexCompilation`. This name stresses the fact that its behaviour is provided by the radial segment class, named `OrganOfCortiComplexSegment`¹². Each `OrganOfCortiComplexSegment` requires some data including inputs, state space models, states, and variables to assign as outputs. The need for separate dynamic and data models was motivated. Inputs might come from the fluid, other segments (via the `read_state()` method), or external inputs such as the neural pathways. Similarly, these sources can also receive the outputs from the segment. The state space models used in this class can be nonlinear. The model of passive dynamics given in (4.33) is used as an example of an instance of an `OrganOfCortiComplexSegment` in figure 4.5. When a method such as `update_states()` in the compilation is called, this method essentially passes on the request, together with any additional data required, to all the segments. Similarly, the `OrganOfCortiComplexSegment` class passes certain requests to the `FluidInterface` class for processing. The basilar membrane (BM) and possibly either the tectorial membrane (TM) or reticular lamina are examples of instances of the `FluidInterface` class.

The `OrganOfCortiComplexCompilation` class also contains a `Parameters` object, which must have some method `calculate_values()` to provide values for the radial segments. Neely (1981a) was one of the first reports investigated in this work, and the basilar membrane admittance function that Neely studied is equivalent to equation (4.33). Therefore, the instance of the `Parameters` object shown in the figure produces the cochlear map from that report. It is also useful to store certain size data concerning the compilation in the class, for example the number segments and scalae.

Figure 4.6 on page 104 shows a similar inheritance diagram for the `OtherBoundary` class. It is simpler for two main reasons: it is not treated as a collection of segments that are able to move independently, so only a single state space model is required; and the class only interfaces with the fluid at one boundary so there is no benefit in creating instances of a separate `FluidInterface` class. In practice, the `OtherBoundary fluid_interface` is more flexible than that in the `OrganOfCortiComplexSegment` because it handles the data mapping functionality of the `OrganOfCortiComplexCompilation` class and manages integration of the input over the boundary. Furthermore, it allows the input to be either pressure/potential or acceleration/velocity (to allow the class to model

¹²`OrganOfCortiComplex` rather than `CochlearPartition` was selected as the prefix for the class name because the models considered when planning the class were best described as models of the organ of Corti complex. Cochlear partition seems to be a slightly more general term, but the functionality of the class is unaffected.

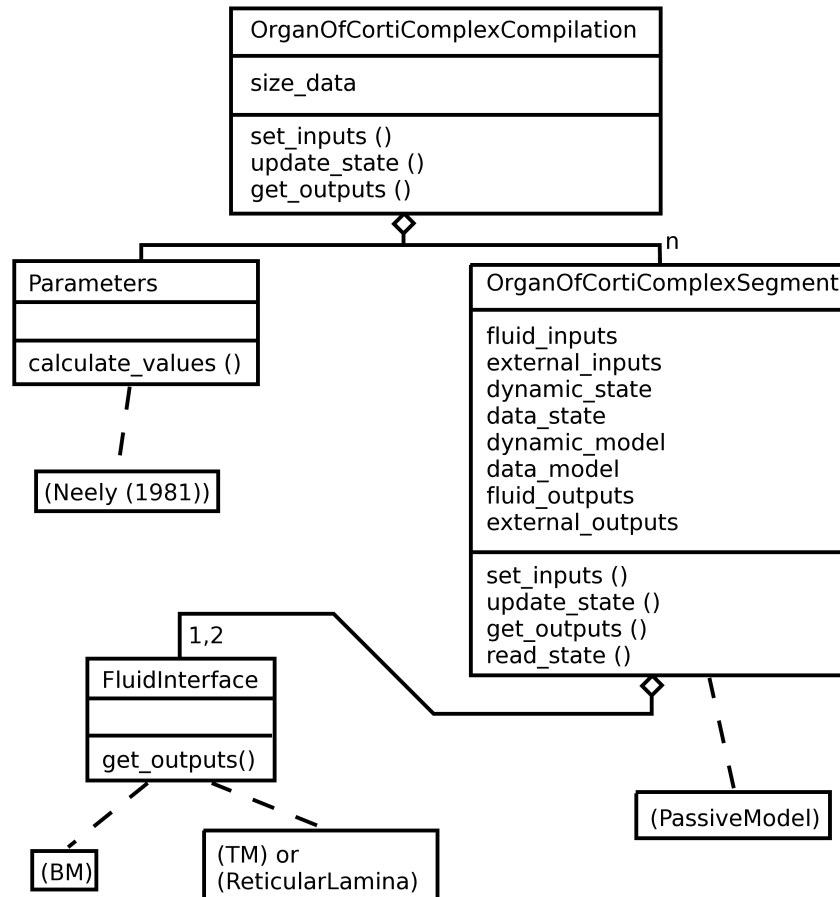


Figure 4.5: An inheritance diagram for the cochlear partition (using `OrganOfCortiComplexCompilation` as its class name). The diagram follows the conventions for the object modelling technique static models (see [Bouzeghoub et al. 1997](#), section 5.6). The entity name, a list of attributes, and a list of operations is provided for each class. A diamond and solid line indicate classes in an aggregation relationship and a number by the line indicates the cardinality. Derived classes are indicated by dashed lines and class names surrounded by brackets. BM, basilar membrane; TM, tectorial membrane.

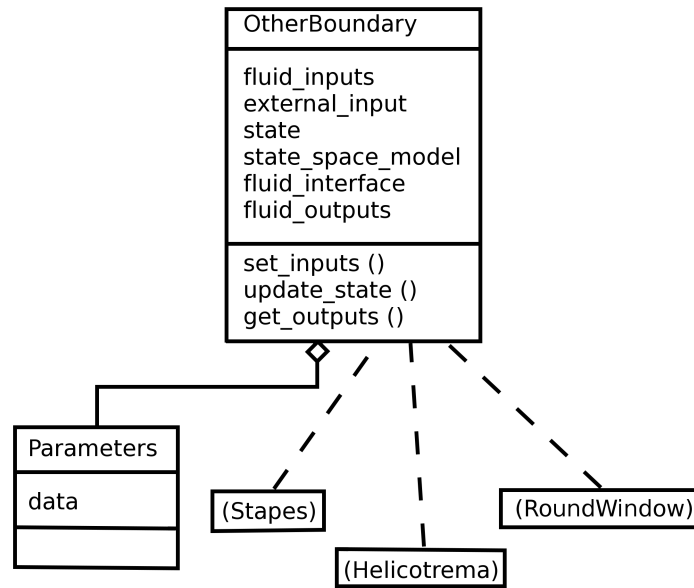


Figure 4.6: An inheritance diagram for the `OtherBoundary` class. Conventions adopted in the diagram are as described in figure 4.5.

the round window boundary conditions), and can handle either scalar or vector fluid formulations. This complexity is hidden from the user who extends the class to create a derived class, such as a `Stapes` class, shown in the figure.

Chapter 5

Fluid-Structure Interaction

In this chapter we tackle the important topic of the fluid-structure interaction in the cochlea, and the restrictions that this places on the proposed modelling framework. I investigated this topic in the course of preparing two publications on the subject. Two summaries of these reports will serve as the bulk of the chapter. Each summary will provide a short introduction to the report. The key results shown in the report will then be presented. Finally, the implications for the framework are discussed. For the interested reader, the publications themselves are reproduced in the appendix. They provide the full background to, and important details about, the results presented. I recommended them to readers with an interest in the process of modelling itself, rather than just the results it can produce. By only providing brief summaries, with the full details available in the papers, I aim to ensure that the thesis remains accessible to a wide audience¹.

The findings from the first paper ([Rapson *et al.* 2012](#)) are summarised in section 5.1, and the paper appears in appendix C. This paper made some brief comments on nonlinearity in cochlear models, and these comments are clarified and extended in section 5.2. The highlights of the second paper ([Rapson and Tapson 2011b](#)) follow in section 5.3. The paper itself appears in appendix D. Finally, the implications of the fluid-structure interaction for the modelling framework are summarised.

The analysis that follows focuses on models that solve for a pressure field in the fluid, but similar results hold after the change of variables required to solve for a potential field. Therefore, the results in this section can also be applied to potential based models.

5.1 The Monolithic Approach

[Allen and Sondhi \(1979\)](#) introduced the first time domain cochlear model. Their report covered a number of new and important developments for the field, including details of a Green's function

¹Section 5.2 and chapter 6 also become quite technical. The reader who only wants a general overview of the ideas can safely omit them.

approach to spatial discretisation, and a consideration of plate equations, as applied to the basilar membrane. However, arguably the most important aspect of the paper, the method used to stabilise the fluid-structure interaction by augmenting the Green’s function kernel, was covered in a single line sentence above equation (28) of that paper². It is interesting that many alternatives to Green’s functions have been used, and cochlear models can easily be assembled without using the plate equations directly, but the idea underlying the augmented kernel that they introduced has, until recently, been the *only* method used to stabilise explicit time stepping algorithms in cochlear models.

Elliott *et al.* (2007) introduced a “state space” approach to cochlear modelling that assembled global system matrices describing the cochlear model. On closer examination, it is evident that their approach is different from, but related to, Allen’s augmentation. Once again, the focus of the paper was not on how the approach achieves stability, but rather on the advantages that having the global state space matrices offers.

A thorough investigation and explanation of how these two methods stabilise the fluid-structure interaction seemed to be long overdue. Furthermore, the fact that there are now two alternatives available for use in cochlear models begs the question of what the relationship between these methods is, when one should be used in preference to the other, and whether either is suitable for the framework proposed. The first paper, reproduced in appendix C, addresses these questions. It argues that both approaches are *monolithic*, in that they either explicitly or implicitly solve a single, large state space matrix representing the dynamics of the system as a whole. Therefore, this section is entitled “the monolithic approach”. The paper also extends both of the methods using the finite element method, which is proposed for this framework. Allen and Sondhi’s approach has been used with the finite element method previously, but the approach proposed by Elliott *et al.* had only been applied to 1D finite difference models.

5.1.1 Key Results

5.1.1.1 Extension of \mathbf{F}^{-1} to 2D using the Finite Element Method

The matrix \mathbf{F}^{-1} in Elliott *et al.* (2007) represents the behaviour of fluid, and it satisfies the equation:

$$\mathbf{p}(t) = \mathbf{F}^{-1}\mathbf{a}(t), \quad (5.1)$$

where \mathbf{F} is a tridiagonal 1D finite difference fluid matrix, $\mathbf{p}(t)$ is a vector of pressures, and $\mathbf{a}(t)$ is a vector of accelerations at the fluid boundary. In Rapson *et al.* (2012), it is shown that \mathbf{F}^{-1} can be

²In a private communication, Jont Allen clarified that he found the technique of augmenting the kernel while studying the eigenvalues of the system.

replaced by an equivalent matrix based on the finite element method. Let:

$$\mathbf{K}\mathbf{p}_n(t) = \mathbf{R}\mathbf{a}(t), \quad (5.2)$$

where \mathbf{K} is the system matrix arising from the Laplace operator in equation (4.14), $\mathbf{p}_n(t)$ is the pressure at the nodes, and \mathbf{R} is a matrix such that $\mathbf{R}\mathbf{a}(t)$ is the integral over the cochlear partition that Gauss quadrature would provide. Furthermore, let a matrix \mathbf{Q} interpolate from the nodal pressure to the pressure at Gauss quadrature points $\mathbf{p}(t)$:

$$\mathbf{p}(t) = \mathbf{Q}\mathbf{p}_n(t). \quad (5.3)$$

Then \mathbf{F}^{-1} can be replaced by:

$$\mathbf{F}^{-1} = \mathbf{Q}\mathbf{K}^{-1}\mathbf{R}, \quad (5.4)$$

where suitable boundary conditions are required on \mathbf{K}^{-1} to allow inversion³. Full details of the finite element spaces required are provided in the article.

5.1.1.2 Expressing Allen and Sondhi (1979) in Terms of Global Matrices

In addition to the matrix \mathbf{F}^{-1} , Elliott *et al.* (2007) defines collected state space matrices representing all the radial segments of the cochlear partition, as well as the model of the stapes boundary condition. These are \mathbf{A}_E , \mathbf{B}_E , and \mathbf{C}_E such that:

$$\dot{\mathbf{X}}(t) = \mathbf{A}_E\mathbf{X}(t) + \mathbf{B}_E\mathbf{p}(t), \quad (5.5)$$

$$\mathbf{a}(t) = \mathbf{C}_E\dot{\mathbf{X}}(t). \quad (5.6)$$

In this case, $\mathbf{X}(t)$ is a vector representing the state of the cochlea as a whole. In particular, it contains the state of each radial segment. The input to the model in that paper is taken as the component of the stapes acceleration due to the middle ear. It is represented as a vector $\mathbf{q}(t)$, with the same length as $\mathbf{a}(t)$.

In section 4.3, the state space output equation was defined to be a function of the state and the input. Equations (5.5) and (5.6) can be expressed in this form. If (5.5) is substituted into (5.6) we obtain:

$$\mathbf{a}(t) = \mathbf{C}_E\mathbf{A}_E\mathbf{X}(t) + \mathbf{C}_E\mathbf{B}_E\mathbf{p}(t), \quad (5.7)$$

$$= \tilde{\mathbf{C}}\mathbf{X}(t) + \tilde{\mathbf{D}}\mathbf{p}(t). \quad (5.8)$$

It turns out that keeping all occurrences of \mathbf{C}_E and \mathbf{B}_E makes the manipulations that are used later to

³Rows of $-\mathbf{F}^{-1}$ are equivalent to the Green's function for specific place ($G(x, \bar{x})$ for fixed x , using the notation of Mammano and Nobili (1993), see especially figure A3 of that paper).

relate [Allen and Sondhi \(1979\)](#) and [Elliott et al. \(2007\)](#) clearer. Therefore, the unusual form of (5.6) is justified.

Using these definitions, it is shown in [Rapson et al. \(2012\)](#) that for the closed loop:

$$\begin{aligned}\dot{\mathbf{X}}(t) &= \left[\mathbf{I} + \mathbf{B}_E \mathbf{Q} [\mathbf{K} - \mathbf{R} \mathbf{C}_E \mathbf{B}_E \mathbf{Q}]^{-1} \mathbf{R} \mathbf{C}_E \right] \mathbf{A}_E \mathbf{X}(t) \\ &\quad + \mathbf{B}_E \mathbf{Q} [\mathbf{K} - \mathbf{R} \mathbf{C}_E \mathbf{B}_E \mathbf{Q}]^{-1} \mathbf{R} \mathbf{q}(t),\end{aligned}\tag{5.9}$$

$$\mathbf{a}(t) = \mathbf{C}_E \dot{\mathbf{X}}(t).\tag{5.10}$$

5.1.1.3 Relationship Between [Allen and Sondhi \(1979\)](#) and [Elliott et al. \(2007\)](#)

[Elliott et al. \(2007\)](#) defined two global state space matrices \mathbf{A} and \mathbf{B} to represent the behaviour of their cochlear model. The overall state equation using these matrices is:

$$\dot{\mathbf{X}}(t) = \mathbf{A} \mathbf{X}(t) + \mathbf{B} \mathbf{F}^{-1} \mathbf{q}(t).\tag{5.11}$$

Comparison between the equations in [Elliott et al. \(2007\)](#) and equation (5.9) shows that:

$$\begin{aligned}\mathbf{A} &= \left[\mathbf{I} - \mathbf{B}_E \mathbf{Q} \mathbf{K}^{-1} \mathbf{R} \mathbf{C}_E \right]^{-1} \mathbf{A}_E \\ &= \left[\mathbf{I} + \mathbf{B}_E \mathbf{Q} [\mathbf{K} - \mathbf{R} \mathbf{C}_E \mathbf{B}_E \mathbf{Q}]^{-1} \mathbf{R} \mathbf{C}_E \right] \mathbf{A}_E,\end{aligned}\tag{5.12}$$

$$\begin{aligned}\mathbf{B} \mathbf{F}^{-1} &= \left[\mathbf{I} - \mathbf{B}_E \mathbf{Q} \mathbf{K}^{-1} \mathbf{R} \mathbf{C}_E \right]^{-1} \mathbf{B}_E \mathbf{Q} \mathbf{K}^{-1} \mathbf{R} \\ &= \mathbf{B}_E \mathbf{Q} [\mathbf{K} - \mathbf{R} \mathbf{C}_E \mathbf{B}_E \mathbf{Q}]^{-1} \mathbf{R},\end{aligned}\tag{5.13}$$

where the first expression in each case is from the [Elliott et al. \(2007\)](#), and the second is the global matrix form of [Allen and Sondhi \(1979\)](#). Equation (5.4) has been used to replace \mathbf{F}^{-1} .

In [Rapson et al. \(2012\)](#), the term *monolithic state space* (MSS) is suggested for methods using the first expressions, and *MSS-variation* (MSSV) for methods using the second expression. Finally, the term *monolithic iterative* (MI) is suggested for methods that perform the computation along the lines of [Allen and Sondhi \(1979\)](#), that is, without forming global matrices explicitly. These methods implicitly solve equation (5.9) in the sequence of steps that they take.

It is shown that an identity for the inverse of a sum of matrices in [Henderson and Searle \(1981\)](#) relates the MSSV and MSS methods. This relationship gives people who design cochlear models the freedom to select the most appropriate starting point for their work, whether for a computational model or a theoretical result (a benefit we will use in chapter 7). That paper also offers other identities and expressions that might also provide useful manipulations of the representation of the system.

5.1.1.4 Performance of the Methods

A number of key performance areas for the methods were assessed. These were: the numerical precision of the methods; the storage requirements in memory; and the number of floating point operations (flops) required.

The fact that the MSS method requires two inversion in equation (5.12), one of which is an inversion of a full matrix, whereas the MSSV method only requires a single inversion, is a recurring feature in the comparison between these methods. The smaller number of inversions required in the MSSV method (furthermore, on better conditioned matrices) means that estimates of the numerical precision and number of flops required favour the MSSV method over the MSS method. As these methods both ultimately form \mathbf{A} , \mathbf{B} , and \mathbf{F}^{-1} , the storage requirements for the methods are similar.

The MI method corresponds to the iterative form of equation (5.9), so the numerical precision properties of MI and MSSV are similar. The memory requirements and flops required are different, however. In the MI method, only the matrix $\mathbf{K} - \mathbf{R}\mathbf{C}_E\mathbf{B}_E\mathbf{Q}$ must be formed and inverted (or rather solved) explicitly. A matrix of a similar size is required by both MSS and MSSV when forming \mathbf{A} , \mathbf{B} , and \mathbf{F}^{-1} . In terms of initialising the model (i.e. the work required before the first time step can be computed), the MI method requires far less computation because the collected matrices and \mathbf{A} , \mathbf{B} , and \mathbf{F}^{-1} are never formed. However, much of the computational effort in a time domain cochlear model is devoted to time steps rather than initialisation. Theoretically, assessments of the asymptotic order of computations required for the methods favours MSSV and MSS over MI. In practice, the limitations on the actual size of cochlear models provided by physiological considerations (chiefly the number of inner hair cells and geometry) means that the full benefit of the asymptotic order is not seen. The calculations performed indicate that the storage requirements and flops needed for a time step in the MI method are comparable to, or better than, the MSS and MSSV methods for all but the largest 3D models.

5.1.1.5 Nonlinear Systems and Implicit Time Stepping Methods

Rapson *et al.* (2012) largely deals with a linear cochlear model, because this allows the key results to be presented clearly and concisely. The discussion considers the extension of the results to nonlinear systems. Researchers from Università di Roma “Tor Vergata” have recently published a number of reports extending Elliott *et al.* (2007) to a range of nonlinear systems, using an implicit time stepping scheme. Certain comments in those papers (in particular in Sisto *et al.* (2010)) seem to imply that implicit time stepping methods are as efficient as explicit methods for *all* nonlinear cochlear models. A number of authors have successfully used explicit methods for nonlinear cochlear models. Depending on which terms in equation (5.12) are affected by the nonlinearity, explicit schemes may have a computational advantage over implicit schemes. \mathbf{A}_E is never inverted, and if the nonlinearity only affects this matrix, then its effects on the performance of an explicit method could be minor.

On the other hand, \mathbf{B}_E and \mathbf{C}_E form part of a matrix that must be solved (inverted) in both expressions for \mathbf{A} . That matrix needs to be inverted each time the linearisation changes, greatly increasing the expense of an explicit method, and making implicit methods more attractive. Section 5.2 will elaborate on these comments.

5.1.2 Implications for the Modelling Framework

5.1.2.1 Do Any of the Methods Suit this Framework?

In this cochlear modelling framework, the emphasis is placed on extensibility and modularity, achieved through the use of computational objects that correspond to physical structures. The MI method is the approach that is most in keeping with this methodology in that the data that characterises each radial segment model (chiefly the state and dynamic model) can be encapsulated into radial segment objects. Forming the global system matrices required by the MSS and MSSV methods goes against this methodology in that state and dynamic model are stored at a global level.

If this modelling framework requires an MI style computation, then we are particularly concerned about the efficiency of that approach. Fortunately, the assessment of the expected performance of the various options confirmed that it is an efficient way to compute cochlear models.

The close correspondence between MI and MSSV is useful for a number of reasons. It means that the framework can use an object orientated approach with an MI style computation, but its behaviour can still be compared to (and verified against) the underlying monolithic system. Furthermore, it provides a pathway to apply results shown for MSS and MSSV to the modelling framework.

5.1.2.2 Fluid and Cochlear Partition Fundamentally Coupled

Equations (5.12) and (5.13) summarise the alternatives available for simulating the system. The MSS style computation involves the inversion of the matrix $\mathbf{I} - \mathbf{B}_E \mathbf{Q} \mathbf{K}^{-1} \mathbf{R} \mathbf{C}_E$, and therefore it is not a good starting point for an object orientated framework as proposed in this work. We noted above that the MI computation is more suitable. However, even for the MI style computation the matrix $\mathbf{K} - \mathbf{R} \mathbf{C}_E \mathbf{B}_E \mathbf{Q}$ must be solved. In this matrix \mathbf{K} , \mathbf{R} , and \mathbf{Q} are related to the fluid model, and \mathbf{C}_E and \mathbf{B}_E are related to the radial segment models. $\mathbf{K} - \mathbf{R} \mathbf{C}_E \mathbf{B}_E \mathbf{Q}$ is the finite element equivalent of Allen's augmented kernel. Therefore, there is a fundamental coupling between the fluid and cochlear partition models that our implementation of the MI method must respect, otherwise instabilities will be introduced.

The effect of this coupling is that the fluid solver must have access to $\mathbf{C}_E \mathbf{B}_E$, as well as the component of the acceleration or velocity that depends on the state of the partition. In the framework, this will be achieved by defining the output of the radial segment models to be a linearisation of the type $a(t) = k_1 p(t) + k_2$ (see section 6.3.1).

5.1.2.3 System Eigenvalues

[Elliott *et al.* \(2007\)](#) stressed the benefits of being able to compute the eigenvalues of a (linearised) cochlear model. Studying the eigenvalues of the \mathbf{A} matrix can quickly show whether the model is linearly stable for a given set of parameters. [Rapson *et al.* \(2012\)](#) reaffirmed the usefulness of eigenvalues. In that research, they were used to identify low frequency whole cochlear resonances in the model, and to visualise the effects of variations in the parameters of the model.

Standard numerical libraries provide robust algorithms for computing the eigenvalues of a matrix, but these require the complete matrix, \mathbf{A} in our case, to be available. Therefore, it is useful to be able to assemble \mathbf{A} when eigenvalues are required, and this requires a slight relaxation of principles of encapsulation (read-only access to the data required for \mathbf{A}_E , \mathbf{B}_E , and \mathbf{C}_E is needed). Approximations to eigenvalues can be found using iterative Krylov space methods, but using these methods to avoid assembling \mathbf{A} may prove to be more complicated, and to require an equivalent relaxation of the data encapsulation. Allowing this relaxation of encapsulation is well within the spirit of the framework proposed: the aim of the framework is to encourage a structured way of assembling cochlear models, without placing unnecessary limitations on the form of the models.

5.1.2.4 The Value of an Extensible Framework

In section 2.1.2 it was argued that that using a single modelling framework to compare different cochlear models can simplify the analysis. The models implemented in [Rapson *et al.* \(2012\)](#) support this claim. In that paper five variations of the same problem were solved. They were: 1D finite difference MSS; 1D finite element MSS; 2D finite element MSS; 2D finite element MSSV; and 2D finite element MI. These codes vary in three aspects: dimension of the problem; spatial discretisation method; and the time stepping method. Of these aspects, only the dimension of the problem should introduce a fundamental change in the solution. With sufficient mesh refinement, and perfect arithmetic, the other factors should not change the solution. In practice, they also have an effect on the solution, however.

Maintaining five codes that should have a close correspondence is not straightforward, as might be expected, but the fact that all the codes could share the same radial segment objects reduced the possibility for errors in the data that the models require. It was easier to locate analysis points for comparisons between models based on the finite element method than between the finite difference model and the finite element models. This was simply because the nodes and quadrature points can be identical in the finite element models, whereas the concepts of nodes and quadrature points are not relevant for a finite difference simulation. Although a close correspondence was obtained for the 1D finite difference MSS model and the 1D finite element MSS model, simply using different method for spatial discretisation of the problem meant that certain points in the domain were easier

to compare than others (particularly where a finite element node and a finite difference grid point coexist). Furthermore, additional post processing was necessary when comparing the results.

Similarly, when comparing MSS codes to other MSS codes, the actual operations and the sequence in which they are applied can be identical. This simplifies the analysis of the error introduced by the sequence of operations, and with other factors being equal, we can expect the solution of two separate codes to be identical. If we compare any pair of dissimilar methods, then the operations and the sequence of operations cannot be the same. Even if we solve the identical model using the two different methods, there will be some difference in their solution. Rapson *et al.* (2012) confirmed that small differences must exist if different methods are used to obtain the solution. These types of differences can confuse the issue when comparing different methods, and by using a single framework they can be avoided.

5.2 More on Nonlinearity

In the context of this chapter on the fluid-structure interaction in the cochlea, it is appropriate to elaborate on the ideas in section 5.1.1.5. We consider a nonlinear segment model that is greatly simplified compared to the flexibility that the specification of the framework allows for the segment model. However, it will allow us to highlight some key implications of nonlinearity. Once again, we will mainly consider a linear fluid model governed by the Laplace equation $\nabla^2 p = 0$, but the results will hold more generally.

5.2.1 The Nonlinear Fluid-Structure Feedback Loop

Section 4.3 provided a general definition for nonlinear state space models. In the case of a state space model representing the cochlear partition in the fluid-structure feedback loop, the input and output variables are clearly defined: they are either $p(t)$ and $a(t)$, or $\phi(t)$ and $v(t)$ (we are assuming a Laplacian fluid model at this point). This allows us to be more specific about the characteristics of the state update equation (4.18) and the output equation (4.19). In this case, the output equation can simply select a state, because $a(t) = \partial v(t)/\partial t$ and $v(t) = \partial d(t)/\partial t$, and the state space models can be expressed with $v(t)$ or $d(t)$ as a state. Therefore, even for a nonlinear model the output can be defined by $\mathbf{C}\dot{\mathbf{X}}_k(t)$, where the matrix \mathbf{C} is not affected by the nonlinearity, and $\dot{\mathbf{X}}_k(t)$ is the state gradient associated with the k^{th} segment, arranged as a vector.

The state gradient is defined by a set of nonlinear state update functions, $\{f_{i,k}\}$, where k indexes the set of functions associated with a specific segment, and i indexes the update equation associated with a particular state. If we use square brackets $[]$ rather than braces $\{\}$, we impose an ordering on

the set, similar to a vector in linear algebra. Using this notation, we can write:

$$\dot{\mathbf{X}}_k(t) = [f_{i,k}(\mathbf{X}_k(t), p_k(t))], \quad (5.14)$$

where $\mathbf{X}_k(t)$ is the state of the k^{th} radial segment, and $p_k(t)$ is its input.

In equation (5.14), we assume that the sets $\{f_{i,k}\}$ are time-invariant and take a single input $p_k(t)$ for simplicity. The main reason to relax the time-invariance constraint would be to allow a mechanism to include neural feedback. This feedback could also be included via the mechanism of providing external inputs. Although external inputs are neither the input nor the output in the fluid-structure feedback loop, they can affect it because $\{f_{i,k}\}$ is nonlinear. In a two scalae cochlear model, the segment would need to accept two inputs, and provide two outputs to account for the scalae. The specification also allowed the fluid input at neighbouring segments to affect the output. Therefore, equation (5.14) involves a number of simplifications with respect to the full flexibility allowed in the specification, but the discussion of this matrix will serve as a useful template for the whole partition model.

To close the feedback loop, we require equations (5.2), (5.14), and (5.3). Once again, $[]$ indicates that a set of values or functions is ordered like a vector. In order to avoid a cumbersome expression, we divide the final equation into its constituent parts to obtain:

$$\mathbf{K}\mathbf{p}_n(t) = \mathbf{R}[\mathbf{C}\dot{\mathbf{X}}_k(t)], \quad (5.15)$$

$$\dot{\mathbf{X}}_k(t) = [f_{i,k}(\mathbf{X}_k(t), p_k(t))], \quad (5.16)$$

$$p_k(t) = (\mathbf{Q}\mathbf{p}_n(t))_k, \quad (5.17)$$

where variables and sets subscripted by k are associated with the k^{th} segment. The expression as a whole is implicit because $\mathbf{p}_n(t)$ appears in both the first and third lines.

If equations (5.15–5.17) can be solved for $\mathbf{p}_n(t)$, then the state gradient for each segment can be calculated from equations (5.16–5.17). Treating the equations as an explicit update expression for $\mathbf{p}_n(t)$, for example by replacing $\mathbf{p}_n(t)$ in (5.15) with $\mathbf{p}_n(t + \Delta t)$, is known to fail in the linear case. Therefore, it is unlikely to be suitable for a nonlinear model. We can either revert to an implicit solver strategy at this stage, or continue to investigate the system further by linearising the segment models.

5.2.2 Linearisation of the State Update Equation

Equation (5.14) can be linearised by taking the Taylor expansion of the expression, provided that the expansion is well-defined around the points of interest. In section 4.3.2, we noted that state space systems already require $\{f_i\}$ to be continuous for a solution to exist, and to satisfy a Lipschitz

condition for the solution to be unique. Unfortunately, these conditions are not sufficient to ensure that $\{f_i\}$ has a Taylor expansion up to linear terms. The major extra requirement is the existence of partial derivatives at the point where the expansion is required.

Consider a specific segment k and time point t , and assume that the conditions for the Taylor expansion are met at that point. Define the state and the input to be \mathbf{X}_t and p_t respectively, then the linearisation of $f_i(\mathbf{X}, p)$ around that point is:

$$f_i(\mathbf{X}, p) \approx f_i(\mathbf{X}_t, p_t) + (\mathbf{X}_j - \mathbf{X}_{t_j}) \frac{\partial f_i}{\partial X_j}(\mathbf{X}_t, p_t) + (p - p_t) \frac{\partial f_i}{\partial p}(\mathbf{X}_t, p_t) \quad (5.18)$$

$$\begin{aligned} &= \left(f_i(\mathbf{X}_t, p_t) - \mathbf{X}_{t_j} \frac{\partial f_i}{\partial X_j}(\mathbf{X}_t, p_t) - p_t \frac{\partial f_i}{\partial p}(\mathbf{X}_t, p_t) \right) \\ &\quad + \left(\frac{\partial f_i}{\partial X_j}(\mathbf{X}_t, p_t) \right) \mathbf{X}_j + \left(\frac{\partial f_i}{\partial p}(\mathbf{X}_t, p_t) \right) p \end{aligned} \quad (5.19)$$

$$= \mathbf{Y}_t + \mathbf{A}_t \mathbf{X} + \mathbf{B}_t p, \quad (5.20)$$

where the time dependence of variables is not explicitly shown, and the index j is used to extract specific states where required. When we linearise the model in this way, we obtain matrices \mathbf{A}_t and \mathbf{B}_t that can be collected to form \mathbf{A}_E and \mathbf{B}_E . In addition, we obtain a new vector \mathbf{Y}_t that represents the constant in the linearisation of $f_i(\mathbf{X}, p)$. (\mathbf{Y}_t is essentially a linear interpolation of $(f_i(\mathbf{X}_t, p_t))$ back to the point $(\mathbf{0}, 0)$. This allows us to deal with the values of $\mathbf{X}(t)$ and $p(t)$ in the following equations rather than differences.) We define \mathbf{Y}_E to be a collected form of \mathbf{Y}_t . Using this linearisation in the expression (5.15–5.17) we obtain:

$$\mathbf{K} \mathbf{p}_n(t) = \mathbf{R} \mathbf{C}_E (\mathbf{Y}_E + \mathbf{A}_E \mathbf{X}(t) + \mathbf{B}_E \mathbf{Q} \mathbf{p}_n(t)), \quad (5.21)$$

where time dependence has been reintroduced. We proceed by collecting $\mathbf{p}_n(t)$ terms:

$$[\mathbf{K} - \mathbf{R} \mathbf{C}_E \mathbf{B}_E \mathbf{Q}] \mathbf{p}_n(t) = \mathbf{R} \mathbf{C}_E \mathbf{Y}_E + \mathbf{R} \mathbf{C}_E \mathbf{A}_E \mathbf{X}(t). \quad (5.22)$$

Using our expression for $\dot{\mathbf{X}}(t)$ again, we finally obtain:

$$\begin{aligned} \dot{\mathbf{X}}(t) &= \mathbf{Y}_E + \mathbf{A}_E \mathbf{X}(t) + \mathbf{B}_E \mathbf{Q} [\mathbf{K} - \mathbf{R} \mathbf{C}_E \mathbf{B}_E \mathbf{Q}]^{-1} \mathbf{R} \mathbf{C}_E (\mathbf{Y}_E + \mathbf{A}_E \mathbf{X}(t)) \\ &= \left[\mathbf{I} + \mathbf{B}_E \mathbf{Q} [\mathbf{K} - \mathbf{R} \mathbf{C}_E \mathbf{B}_E \mathbf{Q}]^{-1} \mathbf{R} \mathbf{C}_E \right] (\mathbf{Y}_E + \mathbf{A}_E \mathbf{X}(t)). \end{aligned} \quad (5.23)$$

This result should be compared to equation (5.9). The two expressions differ only in that the constant terms \mathbf{Y}_E are now added to the contribution of the states $\mathbf{A}_E \mathbf{X}(t)$. (We did not obtain the term multiplied by $\mathbf{q}(t)$ in (5.9) because the input to the cochlear model at the stapes was neglected during the discussion.) Using the identity relating the MSS and MSSV methods, we can obtain the equivalent MSS expression. The important implication of equation (5.23) is that the instantaneous fluid-structure coupling in a nonlinear model is governed by the same expression as for a linear

model. To discuss how that coupling evolves with time, we need to look at what terms affect \mathbf{A}_E , \mathbf{B}_E , and \mathbf{Y}_E .

5.2.3 Dependence of \mathbf{A}_E , \mathbf{B}_E , and \mathbf{Y}_E on $\mathbf{p}(t)$ and $\mathbf{X}(t)$

The pressure variables $\mathbf{p}(t)$ and the state $\mathbf{X}(t)$ are fundamentally different in the state space system representing a cochlear model. By the definition of state space systems, the state is only allowed to evolve as determined by the state gradient. Changes to any of the terms in equation (5.23) change $\dot{\mathbf{X}}(t)$ but *do not* affect $\mathbf{X}(t)$. On the other hand, $\mathbf{p}(t)$ is governed by equation (5.22). Changes to any of the terms in (5.22) change the solution $\mathbf{p}(t)$. Although $\mathbf{p}(t)$ does not appear in equation (5.23), the derivation shows that the equation implicitly depends on the pressure. Using equations (5.19–5.23) we will look at how the system is affected if the linearisation depends on $\mathbf{p}(t)$ and $\mathbf{X}(t)$ separately.

5.2.3.1 Dependence on the State $\mathbf{X}(t)$

\mathbf{A}_E , \mathbf{B}_E , and \mathbf{Y}_E are all related to $\frac{\partial f_i}{\partial X_j}$ and $\frac{\partial f_i}{\partial p}$. If $\frac{\partial f_i}{\partial X_j}$ and $\frac{\partial f_i}{\partial p}$ are functions of $\mathbf{X}(t)$ only, what can be said about \mathbf{A}_E , \mathbf{B}_E , and \mathbf{Y}_E ? In the case of \mathbf{A}_E and \mathbf{B}_E , it is clear that they are also functions of $\mathbf{X}(t)$ only, because they depend only the partials. \mathbf{Y}_E depends on the value of the function $f_i(\mathbf{X}(t), p(t))$, so we must consider this term further.

We are considering the case that $\frac{\partial f_i}{\partial X_j}$ and $\frac{\partial f_i}{\partial p}$ are independent of $p(t)$. This implies that $f_i(\mathbf{X}(t), p(t))$ is at most a linear function of $p(t)$. \mathbf{Y}_t was constructed by mapping $f_i(\mathbf{X}(t), p(t))$ back to the origin using the linearisation at the point (\mathbf{X}_t, p_t) . This method of construction means that any variation due to linear terms is eliminated from \mathbf{Y}_t . Hence, \mathbf{Y}_E is also independent of $\mathbf{p}(t)$ in this case.

\mathbf{A}_E , \mathbf{B}_E , and \mathbf{Y}_E are a valid linearisation of the system as long as the state does not vary too much from the linearisation point, that is $\mathbf{X}(t) \approx \mathbf{X}_t$. In an explicit method, like the Runge-Kutta 4th order method, we require a number of estimates of $\dot{\mathbf{X}}(t)$ for various $\mathbf{X}(t)$. The various $\mathbf{X}(t)$ are likely to be sufficiently different that we must recompute \mathbf{A}_E , \mathbf{B}_E , and \mathbf{Y}_E for each estimate of $\dot{\mathbf{X}}(t)$. A change in \mathbf{A}_E and \mathbf{Y}_E does not affect the term $\left[\mathbf{I} + \mathbf{B}_E \mathbf{Q} [\mathbf{K} - \mathbf{R} \mathbf{C}_E \mathbf{B}_E \mathbf{Q}]^{-1} \mathbf{R} \mathbf{C}_E \right]$ because they do not appear in it. Therefore, this term can be reused even if \mathbf{A}_E and \mathbf{Y}_E change, making it relatively cheap to update these matrices for a change in the state.

On the other hand, if \mathbf{B}_E changes, the term needs be updated, and in particular the expression $[\mathbf{K} - \mathbf{R} \mathbf{C}_E \mathbf{B}_E \mathbf{Q}]^{-1}$ is affected. This expression can be expensive to update if the matrix was inverted explicitly, or if a direct factorisation of the matrix was performed to facilitate computation with the matrix (see [Rapson et al. 2012](#) appendix C.2 for more information about these factorisations). In the case that \mathbf{B}_E changes, the comments in [Sisto et al. \(2010\)](#) that explicit methods may not have a computational advantage over implicit methods become important. If only \mathbf{A}_E and \mathbf{Y}_E are affected by a nonlinearity, then we do not need to linearise the problem at all, because we can

compute the nonlinear equivalent of $\mathbf{A}_E \mathbf{X}(t)$ without linearising the problem, and \mathbf{B}_E can be defined easily.

Are $\frac{\partial f_i}{\partial X_j}$ and $\frac{\partial f_i}{\partial p}$ likely to depend on $\mathbf{X}(t)$ in a cochlear model? For the small movements of the cochlear partition associated with physiologically relevant sound input pressure levels, the passive response⁴ of the cochlea is essentially linear. The inner and outer hair cells are the source of the nonlinearities present in the active response. Most researchers consider deflection of the hair cell stereocilia to be the mechanical input to the hair cells. For the outer hair cells, this deflection is governed by the position of the reticular lamina and the tectorial membrane. In the case of the inner hair cells, it is governed by viscous drag on the stereocilia. Both of these causes of deflection are ultimately tied to the state $\mathbf{X}(t)$. Therefore, $\frac{\partial f_i}{\partial X_j}$ is likely to depend on $\mathbf{X}(t)$ for at least some states.

If deflection of the stereocilia is the effective input to the hair cells, then $\frac{\partial f_i}{\partial p}$ need not be a function of the state of the system, and as we noted this can simplify the computation considerably. However, in [Kim and Xin \(2005\)](#) the feedforward forces that the outer hair cells exert on the partition are treated as an additional nonlinear pressure that is added to the fluid pressure. Whether this model could be manipulated to avoid the dependence of $\frac{\partial f_i}{\partial p}$ on $\mathbf{X}(t)$ or not, it is clear that there are models that can benefit from allowing the dependence in the framework. Alternative theories of outer hair cell function may also make $\frac{\partial f_i}{\partial p}$ dependent on $\mathbf{X}(t)$.

5.2.3.2 Dependence on the Pressure $\mathbf{p}(t)$

The classical view of how fluid pressure affects the state of the partition is based around Newton's second law: the pressure difference across the partition causes a force acting on the partition that contributes towards the acceleration of the partition. This effect is basically linear because the mass of the partition and the surface area exposed to the pressure are unlikely to depend on either the state or the pressure itself. Newton's second law is accounted for by the term $\frac{\partial f_i}{\partial p}$ in the linearisation, which determines \mathbf{B}_E , so \mathbf{B}_E should be independent of $\mathbf{p}(t)$.

How might $\frac{\partial f_i}{\partial X_j}$ be affected by the pressure? If pressure changes in the cochlea were significant enough to modify the physical characteristics of the materials then $\frac{\partial f_i}{\partial X_j}$, and hence \mathbf{A}_E , would be affected, but the pressure fluctuations due to the sound input are too small to achieve this. Therefore, \mathbf{A}_E should also be independent of $\mathbf{p}(t)$ under the classical view of the action of fluid pressure on the partition. This also implies that \mathbf{Y}_E is independent of $\mathbf{p}(t)$ as well. It is good news that \mathbf{B}_E , \mathbf{A}_E , and \mathbf{Y}_E appear to be independent of $\mathbf{p}(t)$: if any of them do depend on $\mathbf{p}(t)$, equation (5.22) remains an implicit function that requires a number of evaluations to approach a solution.

Some researchers propose that the outer hair cells respond to the fluid pressure as the major input. This is a new mechanism that would allow the fluid pressure to affect the state of the system, and we

⁴Section 3.3.2 noted some of the difficulties associated with defining the passive response of the cochlea. To some extent this statement is a result of our definition of active and passive.

must consider its effects on \mathbf{A}_E and \mathbf{B}_E . The outer hair cells show nonlinear behaviour. If $\mathbf{p}(t)$ is the effective input to this nonlinear element, then at least $\frac{\partial f_i}{\partial p}$, and hence \mathbf{B}_E , will be dependent on $\mathbf{p}(t)$. $\frac{\partial f_i}{\partial X_j}$ and \mathbf{A}_E are also likely to be affected, because the outer hair cells contribute to the stiffness of the cochlear partition. As we have noted, if either \mathbf{A}_E or \mathbf{B}_E are dependent on $\mathbf{p}(t)$ then care must be taken in solving equations (5.22–5.23) because they are implicit functions.

Therefore, while dependence of \mathbf{A}_E , \mathbf{B}_E , or \mathbf{Y}_E on $\mathbf{p}(t)$ is not likely to be important in classical cochlear models, it is vital for some of the alternatives to the classical view that have been proposed. This cochlear modelling framework must allow for this dependence in order to be fair to the alternative models.

5.2.4 Implications for the Modelling Framework

This section has shown a possible method for linearising a nonlinear cochlear model using Taylor expansions. The linearisation shows that the fluid-structure coupling for the nonlinear model is closely related to the linear model. This is particularly true if the matrix \mathbf{B}_E is unaffected by the nonlinearities, which is often the case. Studying the linearised model highlights the fact that some nonlinearities require more work to incorporate into a cochlear model than others. The linearisations that depend on the fluid pressure, rather than the cochlear partition state only, are particularly troublesome. However, all the cases discussed can be incorporated into the modelling framework.

5.3 A Partitioned Approach?

[Rapson and Tapson \(2011b\)](#), reproduced in appendix D, studied the behaviour of cochlear models near a Hopf bifurcation. A Hopf bifurcation occurs at a critical value of some parameter(s) of certain systems. Around a (supercritical) Hopf bifurcation, the system response changes smoothly and reversibly from a stable equilibrium to a limit cycle (see [Strogatz 1994](#)). Hopf bifurcation cochlear models (and the closely related Van der Pol oscillator cochlear models) have been widely studied. Our particular interest in these models grew out of our work with Tara Hamilton on silicon cochlea chips ([Hamilton *et al.* 2010](#)).

[Rapson and Tapson \(2011b\)](#) begin by investigating a description in the literature of instability for a model in the vicinity of a Hopf bifurcation point. Given our interest in comparing simulations of Hopf bifurcation cochlear models to the response of silicon cochlea chips around the bifurcation point, any instability in the simulation is of concern to us. The wide use of cochlear models that allow a Hopf bifurcation implies that the instability would also be of general concern.

The paper continues by proposing, and then assessing, a partitioned approach to cochlear modelling. The approaches discussed in section 5.1 above are all monolithic, because they either explicitly or

implicitly combine the fluid and structure dynamics into a single large matrix equation. Partitioned approaches to models with fluid-structure interaction are the opposite of monolithic approaches in the sense that they separate the computation into steps solving the fluid dynamics and steps solving the structure dynamics. Partitioned approaches have certain advantages. Notably, they allow smaller and better conditioned matrices to be solved for each of the relevant domains rather than a single large matrix, and they allow state of the art codes and solvers to be used for each domain. They also have some drawbacks: they tend to be less stable than a monolithic method, and are sometimes more computationally expensive.

In the context of the proposed framework, partitioned solver methods are interesting because of the flexibility they allow to change the equations solved on each domain. In section 4.2.5, it was argued that while the fluid dynamics are well represented by an inviscid, incompressible fluid model, an extensible cochlear modelling framework should allow other fluid models to be used as well. This investigation into a partitioning method is viewed as a step towards that goal. The monolithic equations already allow a wide variety of cochlear partition models to be implemented, so the move towards partitioned methods is not as urgent from the point of view of the organ of Corti complex.

There are two basic types of partitioned methods: sequentially staggered; and iteratively staggered. For a sequentially staggered partitioned method, the fluid and structure domains are evaluated *once* per time step, alternately. In an iteratively staggered partitioned method, the fluid and structure simulations are run *a number of times* each time step until they converge to a solution for that time step, and only then is the next time step started. Iteratively staggered methods are more computationally intensive, but they are required for systems with strong fluid-structure coupling. The cochlea is known to have strong fluid-structure coupling, but in our investigation we aim to avoid the need for an iteratively staggered method by extending the features of Allen's fluid augmentation to create a sequentially staggered method. We investigate sequentially staggered methods first because they are simpler and less computationally intensive, when they are suitable.

5.3.1 Key Results

5.3.1.1 Explicit Time Stepping is Often Adequate

Rapson and Tapson (2011b) showed that the instability described in the literature is not a concern for cochlear models. The original description revolved around the fact that a nonlinear cochlear model discretised using a Green's function was unstable for a value of a parameter (R_{30}) that was thought to be stable. Both theoretical considerations, and an analysis of a small cochlear model using a nonlinear dynamical systems (NLDS) package suggested that the value for R_{30} should be stable.

We found that the value described as unstable for the Green's function model was stable for our finite element model, but it was close to the bifurcation point for the system, so slight differences

between the finite element discretisation and the Green's function discretisation might have placed the original model in the linearly unstable region of the bifurcation. If this is the case, then the unstable response reported would be the desired response of the system.

We also noted that the level of discretisation had an effect on the bifurcation point of the system. For small systems in particular, the boundary conditions also had a significant effect on the bifurcation point. This shows that the theoretical considerations used to determine the expected bifurcation point, which just considered the partition dynamics, can only give a region where bifurcation is possible: the exact bifurcation point is affected by the fluid coupling. The original researchers were unfortunate in that the theoretical value they found happened to correspond to the estimate from the NLDS package for the system they studied, although this does not hold in general.

Time domain simulations for this nonlinear cochlear model were consistent with the behaviour expected in the vicinity of bifurcation points, which were located using the eigenvalues of the various systems studied. The finding that explicit time stepping is adequate for this system agrees with the discussion in sections 5.1.1.5 and 5.2, because the nonlinearities studied only affect \mathbf{A}_E .

5.3.1.2 Fast and Slow Interaction Offers Opportunity for Partitioning

The function of the augmentation of the cochlea fluid introduced in [Allen and Sondhi \(1979\)](#) can be viewed in terms of fast and slow feedback paths. Consider equation (5.7):

$$\mathbf{a}(t) = \mathbf{C}_E \mathbf{A}_E \mathbf{X}(t) + \mathbf{C}_E \mathbf{B}_E \mathbf{p}(t). \quad (5.24)$$

The term $\mathbf{C}_E \mathbf{A}_E$ multiplies the state of the system, which is restricted to varying as $\mathbf{X}(t + \Delta t) = \mathbf{X}(t) + \Delta t \dot{\mathbf{X}}(t)$. This implies that as $\Delta t \rightarrow 0$, $\mathbf{X}(t + \Delta t) = \mathbf{X}(t)$. In this sense, the term $\mathbf{C}_E \mathbf{A}_E \mathbf{X}(t)$ varies over “slow” times scales. By contrast, the term $\mathbf{C}_E \mathbf{B}_E$ multiplies the pressure $\mathbf{p}(t)$, and any change in the $\mathbf{p}(t)$ has an immediate effect on $\mathbf{a}(t)$. If the fluid dynamics are specified by equation (4.14) then any change in the acceleration also has an immediate effect on the pressure via an equation of the type $\mathbf{p}(t) = \mathbf{F}^{-1} \mathbf{a}(t)$ (see section 5.1.1.1). Therefore, changes in either the pressure or acceleration have an immediate effect on the other variable via the matrices $\mathbf{C}_E \mathbf{B}_E$ and \mathbf{F}^{-1} , so the feedback via $\mathbf{C}_E \mathbf{B}_E$ occurs over “fast” time scales. Allen's fluid augmentation incorporates the “fast” feedback from the cochlear partition into the fluid model, so that it is available to the fluid solver, and neglects the “slow” feedback path that can be accounted for by updating $\mathbf{X}(t)$ when solutions to the fluid domain are available.

As the fast feedback path through the structure is already accounted for using the fluid augmentation, it would appear that we are halfway to a sequentially partitioned cochlear model. If the fast feedback path through the fluid was accounted for in the structure (cochlear partition) model, then the slower feedback pathways should be adequately modelled by passing values for each step between the domains.

The question, then, is whether we can partition the fluid model into a “fast” component and a “slow” component. The short answer is no: the Laplace equation itself is already the “fast” component of the linear Navier-Stokes equation⁵, and it has no time dependent terms itself. However, during the investigation into the reported instability, it was noted that \mathbf{F}^{-1} has a strongly dominant diagonal in a two segment cochlear model. If this feature extended to larger cochlear models it would offer an alternative to partitioning the fluid model along the lines of “fast” and “slow” feedback, with “local” and “nonlocal” feedback being the logical replacement. The results for the small system showed promise, but a strict “local” versus “nonlocal” partitioning performed poorly on larger systems. Including the effects of the neighbouring segments improved the response considerably, and this can be done using the provisions made for physical coupling between segments, but it seems to be an abuse of the mechanism.

If the Laplace equation is adequate for the cochlear model, there is no reason to use a partitioning scheme along these lines. Partitioning along the lines of “fast” and “slow” components for the Navier-Stokes equations seems to be possible based on the arguments set out, but this work has not tested the idea further. In a cochlear model using a Navier-Stokes fluid model, this could replace some evaluations of the fluid with the Laplace equation that is cheaper to compute. There will certainly be complexities related to the fact that the Laplace equation uses a scalar formulation whereas the Navier-Stokes equations use a vector formulation.

5.3.1.3 Interaction Between Segments is Important

The partitioning into “local” and “nonlocal” components described above ultimately failed because the behaviour of a particular radial segment is significantly affected by the fluid feedback from other segments in the cochlea. The results described by [Rapson and Tapson \(2011b\)](#) show the beginnings of an analysis of the coupling between radial segments through the fluid that will be more fully developed in chapter 7. Based on both the performance of the system where coupling between 6 neighbouring segments is maintained, and plots of the fluid coupling matrix $(\mathbf{Q}[\mathbf{K} - \mathbf{R}\mathbf{C}_E\mathbf{B}_E\mathbf{Q}]^{-1}\mathbf{R})$, it is apparent that segments are more sensitive to their neighbours than the remainder of the segments. This will have important ramifications when we consider whether the fluid is able to provide the inter-segment coupling required for a true travelling wave to exist in the cochlea.

5.3.2 Implications for the Modelling Framework

The investigation into the results of [Matthews and Molnar \(1989\)](#) has again highlighted some of the points raised in section above. These are:

⁵For rapid fluctuations (essentially high frequencies) the inertia of the fluid has a larger effect than the viscosity (see section 4.2.3).

- Explicit time stepping can be adequate for nonlinear cochlear models;
- The eigenvalues of a cochlear model can provide important information about its performance, justifying an exception to strict encapsulation of data to allow their computation;
- The response of the same model in different implementations should be expected to be subtly different. In the vicinity of a critical point of a system, these subtle differences can have a significant effect on the response of that system. This supports one of the central claims in this thesis: that the ability to simulate a wide variety of models in the same framework will be beneficial to the cochlear modelling community.

The investigation into a sequentially staggered partitioning method under predetermined constraints was unsuccessful as far as it was taken. The choice of a sequentially staggered scheme rather than an iteratively staggered scheme was based on the desire to investigate simple solutions before complex, as well as the knowledge that if a partitioned method is to replace the monolithic system based on the Laplace equation then it must be computationally cheap to implement.

Partitioning methods will show their value when fluid models other than the Laplace model are required. Therefore, the concept of partitioning along the lines of “fast” and “slow” components in the response has not been given a fair test yet, and I think that it still has promise. This work also shows that interaction between segments via the fluid is important for accurate simulation of cochlear models. The implementation of the fluid-structure interaction in the framework must correctly capture these interactions. It is particularly important that the fast feedback loop between the fluid and the cochlear partition over small neighbourhoods is well represented. The Laplace model is a good model for interactions over short time scales and small distances, so it characterises this feedback loop well.

5.4 Summary

In this section, two papers that discuss the interaction of the fluid and structure models have been overviewed. A number of important points have been raised, some of which lend support for the goals of this thesis, but others place constraints on the way that the specification of the framework from chapter 4 is implemented. A method of linearising nonlinear cochlear models has been presented, and various scenarios that can arise with different nonlinear models have been discussed.

In the two papers, it was apparent that direct comparisons between the same model implemented in different ways is difficult. This serves as motivation for a single framework that can implement different models. We have noted that a specific framework will not necessarily be equally fair to all models, and as the thesis develops it is becoming increasingly apparent that every framework can

at best implement a subset of the cochlear models that might be proposed, but this does not prevent work on a framework that can implement as large a set of models as fairly as possible.

Consideration of the fluid-structure interaction has clarified the constraints it places on the model. In chapter 4 we only noted that these constraints would be important. The major constraints are:

- The behaviour of the cochlear model must correspond to the behaviour of the underlying monolithic system. For the Laplace fluid model with a linear partition model this behaviour is defined by equations (5.11–5.13);
- The MI method is compatible with the goals of the framework (for the Laplace fluid model). This method requires the output from the segment models to be a linearisation based on the fluid domain solution variable, rather than a specific value;
- The segment objects should provide read-only access to data that might be required for the eigenvalues of the monolithic system to be computed.

Maintaining close correspondence between the framework and the monolithic system equations provides an important means to verify the implementation as well as a natural pathway for theoretical results to be applied to implemented models.

The extensibility of our proposed framework rests on the claim that the fluid-structure interaction has a characteristic form. This chapter has supported that claim in a number of ways:

- Equation (5.12) defines the characteristic form in a passive model;
- Equation (5.23) shows that active models share the characteristic form seen in passive models; and
- Section 5.3 showed that the Laplacian fluid model is a particularly good representative of fluid as far as the fluid-structure coupling is concerned.

Chapter 6

Detailed Description of the Framework

6.1 Introduction

We began with a summary of who the users of cochlear models are, and what use they make of them. We asked the question of whether a computational framework can support the full spectrum of possible models. The thesis continued by noting the properties of the cochlea, and numerical tools available that either permit desirable simplifications of the model, or force limitations in the variety of models that can be considered. The discussion has necessitated a gradual shift from general and abstract to specific but concrete. In this chapter, we complete the shift by providing the fine details required for models that use the Laplace fluid model. This important fluid model serves as an example of how other fluid models would be implemented. The chapter follows a similar outline to chapter 4, and develops the relevant sections in that chapter further. Once again, we use the case of a pressure fluid formulation in our expressions, with the understanding that all the expressions have close parallels for the case of a potential fluid formulation.

6.2 The Fluid Model

In this section, we show details of the finite element formulation required for the framework. The implications of chapter 5 on the fluid-structure interaction in the cochlea are incorporated into the formulation, and some conventions for the object interface are selected. For generality, the discussion assumes that the scala vestibuli and the scala tympani are modelled separately.

We first review the conventions that are required for the framework. Other options are possible for all the conventions, but conformance with the proposed conventions is suggested to avoid confusion when reusing objects in new models. We then work through the finite element details for the model. Finally, we note the relationship between the approach presented in this chapter and the monolithic equations discussed in chapter 5.

6.2.1 Sign Conventions

In section 4.2.4, expressions for the pressure and potential Laplace fluid equations were provided. We now select sign conventions that are suitable for those equations. In the case of a pressure field, physics provides a relationship between the acceleration and the gradient of the pressure field, $\mathbf{a} = -\nabla p/\rho$, which serves as a natural choice for one sign convention. For the potential formulation, a similar relationship is adopted, and we follow [Lyon and Mead \(1988b\)](#) by defining $\mathbf{v} = -\nabla\phi$. This definition is convenient because it means that the two expressions only differ by a positive constant, the fluid density ρ . We must still specify the direction of positive velocity and acceleration. For simple cochlear geometries, a Cartesian coordinate system is convenient, and components in the direction of increasing x and y are chosen to be positive. Other systems might be more convenient for different geometries. In particular, polar or curvilinear coordinate systems can be useful when studying coiling in the cochlea (for example, see [Cai et al. 2005](#)).

6.2.2 Segment Input/Output Conventions

The major input to the organ of Corti segment is the pressure (or potential) in the scala vestibuli and the scala tympani. For generality, these must be provided as separate values. Therefore, the segment input is defined as: $\mathbf{p}_{occ} = \begin{bmatrix} p_{sv} & p_{st} \end{bmatrix}^T$, where the subscript sv stands for the scala vestibuli (with scala media), and st stands for scala tympani.

The output from the segment is the acceleration (or velocity) of the membranes in contact with the fluid. This choice is determined by the boundary condition requirements of the fluid model. Furthermore, only the motion *perpendicular* to the boundary can be specified in this type of fluid model. The positive direction is specified such that both Δa_{sv} and Δa_{st} are positive for positive Δp_{sv} , where Δ indicates a small change in the variable. $\bar{\mathbf{n}}_{occ}$ is defined to be a unit normal in this direction. This results in the convention that positive acceleration of the partition is away from the scala vestibuli and towards the scala tympani.

The strong fluid-structure coupling in the cochlea requires the organ of Corti segments to provide a linear expression for acceleration in terms of the fluid pressure, rather than a single number as feedback. This allows the “ $\mathbf{C}_E \mathbf{B}_E$ ” values (related to the segment mass) to be incorporated into the fluid model as required (see section 5.1.2). Expressions of the type $a_{sv} = \mathbf{k}_{1sv} \mathbf{p}_{occ} + k_{2sv}$, where $\mathbf{k}_1 = \begin{bmatrix} k_{10} & k_{11} \end{bmatrix}$ is a row vector of constants, and k_2 is another constant, are sufficient to satisfy the fluid-structure coupling requirements. In general, \mathbf{k}_1 and k_2 must be re-evaluated at each time instant.

6.2.3 Weakform Derivation

The model calls for the following strong form differential equations to be solved on the fluid domain:

$$\nabla^2 p = 0 \quad \text{on } \Omega, \quad (6.1)$$

$$\nabla p \cdot \bar{n} = -\rho a_{sv} \quad \text{on } \Gamma_{sv}, \quad (6.2)$$

$$\nabla p \cdot \bar{n} = -\rho a_{st} \quad \text{on } \Gamma_{st}, \quad (6.3)$$

$$\nabla p \cdot \bar{n} = -\rho a_s \quad \text{on } \Gamma_s, \quad (6.4)$$

$$p = 0 \quad \text{on } \Gamma_{rw}, \quad (6.5)$$

where \bar{n} is a unit outward normal, Ω is the fluid domain, and Γ_{sv} , Γ_{st} , Γ_s , and Γ_{rw} refer to the scala vestibuli, scala tympani, stapes, and round window boundaries respectively. We assume that a zero pressure boundary condition is imposed at the round window for simplicity. The weak form is found by first multiplying the equation on the domain by a trial function, w , and integrating over the fluid domain, to give:

$$\int_{\Omega} w \nabla^2 p d\Omega = 0 \quad \forall w,$$

where \forall means “for all”. Then Green’s theorem is applied, leading to the following equation:

$$\int_{\Omega} \nabla w \cdot \nabla p d\Omega - \oint_{\Gamma} w \nabla p \cdot \bar{n} d\Gamma = 0 \quad \forall w.$$

Substituting for ∇p in the second term leads to the desired weak form:

$$\int_{\Omega} \nabla w \cdot \nabla p d\Omega + \rho \oint_{\Gamma} w \mathbf{a} \cdot \bar{n} d\Gamma = 0 \quad \forall w. \quad (6.6)$$

The first term in this equation is a standard form in the finite element method, and will not be discussed further. The remainder of this section will discuss the second term, because it must be developed further to fit the specific boundary conditions required for the modelling framework.

6.2.3.1 Organ of Corti Complex Boundary Conditions

Figure 6.1 on the following page shows the two fluid bodies that are in contact with the organ of Corti: the scala vestibuli, with cell A a representative of the finite element cells on the boundary; and the scala tympani, with cell C as a similar representative. The definition of positive motion for the segment is indicated by \bar{n}_{occ} in figure. Building the term $\oint_{\Gamma} w \mathbf{a} \cdot \bar{n} d\Gamma$ on Γ_{sv} and Γ_{st} requires that the sign conventions described in previous sections are applied, and that the linearisation for the accelerations a_{sv} and a_{st} respectively are substituted.

In the case of the scala vestibuli (cell A) the component of \mathbf{a} in the direction of the outward normal \bar{n}_A is just a_{sv} , thus the term to be built on cell A is $+\rho \int_{\Gamma_{sv}} w a_{sv} d\Gamma_{sv}$. The outward normal \bar{n}_C is

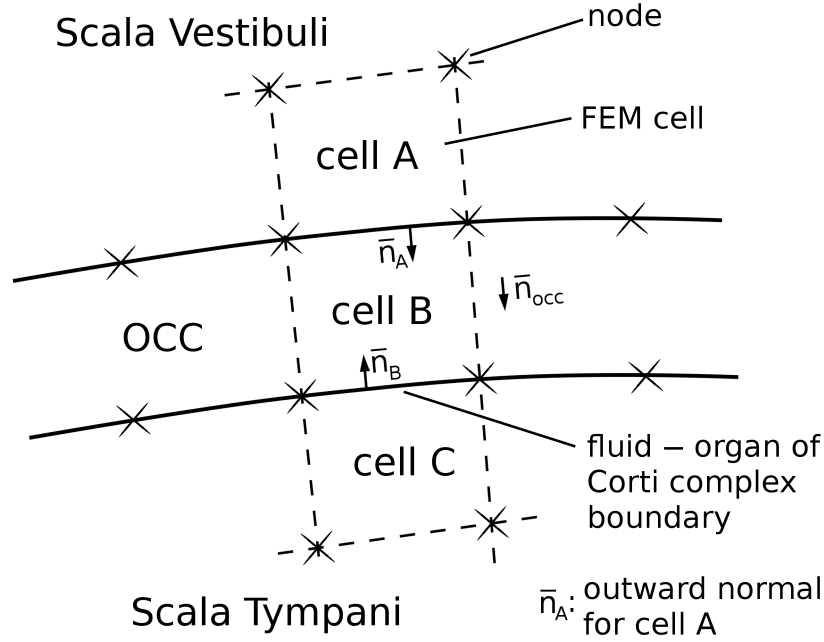


Figure 6.1: Diagram showing finite element cells beside the organ of Corti complex (OCC) in the scala vestibuli and scala tympani (cells A and C respectively). The phantom cell B, used to simplify the process of finding adjacent cells, is also shown. Important normals are indicated.

antiparallel to the definition of a_{st} , hence for the scala tympani the term to be built on cell C is $-\rho \int_{\Gamma_{st}} w a_{st} d\Gamma_{st}$.

In practice, it is convenient to introduce cell B, which does not exist in the fluid region. These interior cells are used to ensure that the cells in the scala vestibuli and scala tympani that are in contact with a give radial segment can be located easily when assigning segments and extracting data from the fluid model. The line integral over the boundary between the fluid and organ of Corti complex can be performed on cell B instead of cells A or C, because the boundaries are shared by cell B. The terms discussed above are valid for integration on either cell, because the outward normals have already been considered. Substituting for the acceleration expressions, the contribution to be assembled on cell B are:

$$\int_{\Gamma_{sv+st}} w \mathbf{a} \cdot \bar{\mathbf{n}} d\Gamma_{sv+st} = \int_{\Gamma_{sv}} w \mathbf{k}_{1sv} \mathbf{p}_{occ} d\Gamma_{sv} + \int_{\Gamma_{sv}} w k_{2sv} d\Gamma_{sv} - \int_{\Gamma_{st}} w \mathbf{k}_{1st} \mathbf{p}_{occ} d\Gamma_{st} - \int_{\Gamma_{st}} w k_{2st} d\Gamma_{st}. \quad (6.7)$$

This expression can be substituted back into (6.6) to obtain:

$$\int_{\Omega} \nabla w \cdot \nabla p d\Omega + \rho \int_{\Gamma_{sv}} w \mathbf{k}_{1sv} \mathbf{p}_{occ} d\Gamma_{sv} - \rho \int_{\Gamma_{st}} w \mathbf{k}_{1st} \mathbf{p}_{occ} d\Gamma_{st} = -\rho \int_{\Gamma_{sv}} w k_{2sv} d\Gamma_{sv} + \rho \int_{\Gamma_{st}} w k_{2st} d\Gamma_{st} - \rho \int_{\Gamma_o} w \mathbf{a} \cdot \bar{\mathbf{n}} d\Gamma_o \quad \forall w, \quad (6.8)$$

where terms dependent on the fluid solution (p) have been collected to the left-hand side. The contribution from the other boundaries, $\int_{\Gamma_o} w \mathbf{a} \cdot \bar{\mathbf{n}} d\Gamma_o$, will be discussed in the next section.

In equation (6.8), all the values required for \mathbf{p}_{occ} are available on cell B. This means that locations of nonzero terms in the matrix assembled from equation (6.8), the “sparsity pattern” for the problem, are defined by the finite element mesh only. If a_{sv} and a_{st} depend on the values that do not lie on cell B, as can be the case for feedforward mechanisms, then extra nonzero terms must be added to the sparsity pattern. The contribution from these mechanisms would also appear on the left-hand side of the expression with the other pressure unknowns.

6.2.3.2 Other Boundary Conditions

So far we have considered the scala vestibuli and scala tympani boundary conditions. The term $\int_{\Gamma_o} w \mathbf{a} \cdot \bar{\mathbf{n}} d\Gamma_o$ in equation (6.8) accounts for the other boundary conditions on the model, and these are discussed further here. Sections 4.5 and 4.6 have already discussed the relationship between the various boundaries and the organ of Corti complex boundary in some detail. There are three important cases to consider:

- When $\mathbf{a} \cdot \bar{\mathbf{n}}$ on boundary Γ_{o_1} is specified as a value a_{o_1} into the region at a each time instant, the integral $-\int_{\Gamma_{o_1}} w a_{o_1} d\Gamma_{o_1}$ describes the boundary, and appears on the right-hand side of (6.8).
- If the boundary requires an expression for $\mathbf{a} \cdot \bar{\mathbf{n}}$, such as $a_{o_1} = k_1 p_{o_1} + k_2$, then the contribution from the boundary will divide between the left- and right-hand sides of (6.8) in a similar fashion to the organ of Corti complex contribution. For the organ of Corti complex, the expression for $\mathbf{a} \cdot \bar{\mathbf{n}}$ must be used to ensure stability of the model, but for the other boundaries stability is generally not at stake. Instead, expressions are required to ensure that the underlying monolithic system is correctly captured.
- Finally, the pressure might be imposed at the boundary, as in equation (6.5). This type of boundary condition is imposed directly by specifying the space of trial solutions to consider for p , and it does not appear in equation (6.8).

In section 4.5.1, we noted that the stapes can be modelled as a rigid body. If this is done then the *integral* of the pressure appears in the expression for $\mathbf{a} \cdot \bar{\mathbf{n}}$ from point 2. We obtain:

$$\int_{\Gamma_s} w \mathbf{a} \cdot \bar{\mathbf{n}} d\Gamma_s = - \int_{\Gamma_s} w k_1 \int_{\Gamma_s} p d\Gamma_s d\Gamma_s - \int_{\Gamma_s} w k_2 d\Gamma_s. \quad (6.9)$$

The double integral over Γ_s is not problematic to implement, but does require a contribution from all nodes on the stapes, which means that extra nonzero terms need to be added to the sparsity pattern.

6.2.4 Function Space Details

So far, we have discussed the various elements that are required for the fluid model in a slightly informal manner. We conclude the section by providing a full statement of the weak form, introducing approximation spaces, and giving a statement of the finite element mass matrix and right-hand side required for the problem. (We follow the notation used in [Fish and Belytschko \(2007\)](#) in this section, and the reader is directed to that text for further background on the finite element method.)

Let $U = \{p | p \in H^1, p = 0 \text{ on } \Gamma_{rw}\}$ and $U_0 = \{w | w \in H^1, w = 0 \text{ on } \Gamma_{rw}\}$ be the functional analysis spaces for the trial solution and the weighting function respectively, where H^1 is the functional analysis space in which both the function and its derivative are square integrable. Using these spaces the weak form is:

Find $p \in U$ such that:

$$\begin{aligned} \int_{\Omega} (\nabla w)^T \nabla p \, d\Omega + \rho \int_{\Gamma_{sv}} w^T \mathbf{k}_{1sv} \mathbf{p}_{occ} \, d\Gamma_{sv} - \rho \int_{\Gamma_{st}} w^T \mathbf{k}_{1st} \mathbf{p}_{occ} \, d\Gamma_{st} = \\ -\rho \int_{\Gamma_{sv}} w^T k_{2sv} \, d\Gamma_{sv} + \rho \int_{\Gamma_{st}} w^T k_{2st} \, d\Gamma_{st} - \rho \int_{\Gamma_o} w^T \mathbf{a} \cdot \bar{\mathbf{n}} \, d\Gamma_o \quad \forall w, \end{aligned} \quad (6.10)$$

where T is the transpose operator, which will allow us to write inner products as matrix multiplications.

Let the continuous spaces be replaced by discrete spaces such that:

$$p(x, y, t) \approx p^h(x, y, t) = \mathbf{N}(x, y) \mathbf{p}_n(t), \quad (6.11)$$

$$w(x, y, t) \approx w^h(x, y, t) = \mathbf{N}(x, y) \mathbf{w}_n(t), \quad (6.12)$$

$$\nabla p(x, y, t) \approx \nabla p^h(x, y, t) = \mathbf{B}(x, y) \mathbf{p}_n(t), \quad (6.13)$$

$$\nabla \omega(x, y, t) \approx \nabla \omega^h(x, y, t) = \mathbf{B}(x, y) \mathbf{w}_n(t), \quad (6.14)$$

where $\mathbf{N}(x, y) \in H^1$ are shape functions, and $\mathbf{B}(x, y)$ are gradients of shape functions. In this case we explicitly show dependence on location and time instant. These are both formally $1 \times n$ vectors, where n is the number of finite element nodes in the domain, although they will only be nonzero for nodes in the finite element cell around the point (x, y) . $\mathbf{p}_n(t)$ and $\mathbf{w}_n(t)$ are the nodal pressures and weighting functions respectively, $n \times 1$.

Finally, let $\mathbf{w}_F(t)$ be the subset of values in $\mathbf{w}_n(t)$ not on Γ_{rw} (on Γ_{rw} , the value of $\mathbf{p}_n(t)$ is enforced separately). Then the weak form becomes:

$$\mathbf{w}_n^T(t) \mathbf{K}_{aug} \mathbf{p}_n(t) = \mathbf{w}_n^T(t) \mathbf{R}_{aug} \quad \forall \mathbf{w}_F, \quad (6.15)$$

where:

$$\begin{aligned}
\mathbf{K}_{aug} &= \int_{\Omega} \mathbf{B}^T(x,y)\mathbf{B}(x,y) d\Omega \\
&+ \rho \int_{\Gamma_{sv}} \mathbf{N}^T(x,y)\mathbf{k}_{1sv0}\mathbf{N}(x,y) d\Gamma_{sv} \\
&+ \rho \int_{\Gamma_{sv}} \mathbf{N}^T(x,y)\mathbf{k}_{1sv1}\mathbf{N}(x_{st},y_{st}) d\Gamma_{sv} \\
&- \rho \int_{\Gamma_{st}} \mathbf{N}^T(x,y)\mathbf{k}_{1st0}\mathbf{N}(x_{sv},y_{sv}) d\Gamma_{st} \\
&- \rho \int_{\Gamma_{st}} \mathbf{N}^T(x,y)\mathbf{k}_{1st1}\mathbf{N}(x,y) d\Gamma_{st}, \tag{6.16}
\end{aligned}$$

$$\begin{aligned}
\mathbf{R}_{aug} &= -\rho \int_{\Gamma_{sv}} \mathbf{N}^T(x,y)k_{2sv} d\Gamma_{sv} + \rho \int_{\Gamma_{st}} \mathbf{N}^T(x,y)k_{2st} d\Gamma_{st} \\
&- \rho \int_{\Gamma_o} \mathbf{N}^T(x,y)\mathbf{a} \cdot \bar{\mathbf{n}} d\Gamma_o. \tag{6.17}
\end{aligned}$$

\mathbf{K}_{arg} is an $n \times n$ matrix and \mathbf{R}_{aug} is a $1 \times n$ vector. The term \mathbf{p}_{occ} requires shape functions evaluated at the point on the current boundary $\mathbf{N}(x,y)$, as well as the opposite point on the other face of cell B, either $\mathbf{N}(x_{st},y_{st})$ or $\mathbf{N}(x_{sv},y_{sv})$. We use indices to extract the correct components of \mathbf{k}_{1sv} and \mathbf{k}_{1st} for multiplication by each term of \mathbf{p}_{occ} . The integration can be performed using Gauss quadrature. k_{2sv} and k_{2st} are required at quadrature points, so it is convenient to define the radial segments to be located at these quadrature points. The arbitrary weighting function $\mathbf{w}_n^T(t)$ in (6.15) has now served its purpose, and can be cancelled. Equation (6.15) is implemented in deal.II using the shape functions that library provides.

This completes the detailed description of the fluid model and its finite element implementation, but we take the opportunity to note the relationship between the expressions developed and the monolithic system. \mathbf{K}_{aug} and \mathbf{R}_{aug} can be related back to the expressions in section 5.1.1 as follows:

$$\mathbf{K}_{aug} = \mathbf{K} - \mathbf{RC}_E\mathbf{B}_E\mathbf{Q}, \tag{6.18}$$

$$\mathbf{R}_{aug} = \mathbf{RC}_E\mathbf{A}_E\mathbf{X}(t) + \mathbf{Rq}(t). \tag{6.19}$$

$\mathbf{C}_E\mathbf{B}_E$ is related to \mathbf{k}_{1sv} and \mathbf{k}_{1st} , and $\mathbf{C}_E\mathbf{A}_E\mathbf{X}(t)$ is related to k_{2sv} and k_{2st} . This close relationship ensures that the monolithic system is correctly implemented in the framework.

6.3 Focusing on the Organ of Corti Complex

In this section, we focus on the `OrganOfCortiComplexSegment` and `OrganOfCortiComplexCompilation` classes, and describe how they are implemented in further detail. We begin by discussing the main feedback loops in the cochlear model.

6.3.1 Fluid Inputs to Output Linearisation

6.3.1.1 State Update Equation

Section 5.2 considered the feedback loop from the the segment's local fluid input to its output under certain simplifications. Notably, it was assumed that the segment output was the gradient of the first state of the dynamic model. A Taylor expansion was used to linearise the state update equation at a time instant t . The discussion that followed described the dependence of the linearisation on the state of the system and the fluid input. It also noted the conditions that allow the linearisation process to be avoided altogether.

The framework provides two mechanisms to allow nonlinearities to be incorporated into the dynamic model. Firstly, a method `update_matrices` is provided in the `OrganOfCortiComplexCompilation` class. A call to this method is passed on to the underlying `OrganOfCortiComplexSegments`, which in turn updates the linearisation of the dynamic model in some predefined manner. (The method is marked as `virtual` in the `OrganOfCortiComplexSegment` class so that it can be replaced as desired in derived classes.) Section 5.2 showed that the nature of the nonlinearities in the dynamic model determine how often the linearisation must be recomputed. Providing a separate function that can be called only when required avoids unnecessary work computing the linearisation.

The second method involves the implementation of the state gradient function itself: the method `dynamic_states_gradient_estimate` that provides $\dot{\mathbf{X}}$ to the `update_states` method in the `OrganOfCortiComplexSegment` class is marked as `virtual`. This allows the linear model provided by the segment class to be replaced by a nonlinear function in implementations. Either, or even a combination of the two methods can be used to implement a nonlinear model.

6.3.1.2 Output Equation

In order to satisfy the fluid-structure coupling requirements, we need to provide output from the organ of Corti complex segments in the form of a linearisation of the type $a_{sv} = \mathbf{k}_{1sv}\mathbf{p}_{occ} + k_{2sv}$. A Taylor expansion is ideal for computing this linearisation because it has the desirable property that a small change in the pressure, $\Delta\mathbf{p}_{occ}$, will produce a change Δa_{sv} in the feedback variable that has the same sign as a similar change in the original nonlinear function. This should facilitate convergence of the linear solver in the fluid model, but the property is not guaranteed for some linearisation techniques. (For example, if we chose to linearise the function $a = v^4 - v^2$ at the point $v_t = 0.75$ as $a = (v_t^3 - v_t)v$, we obtain the linearisation $a = -0.3281v$. The gradient at this point is 0.1875, so although the linearisation might give a fair approximation to a for $|v - v_t| \ll 1$, an error $\Delta\mathbf{p}_{occ}$ will produce a correction Δa_{sv} of the wrong sign, and this might lead to instability. If a Taylor expansion at the point exists, and it is used, at least the sense of the feedback will be correct.)

The linearisation process follows along the lines of (5.18–5.19):

$$g_{sv,i}(\mathbf{X}, \mathbf{p}_{occ}) \approx \left(g_i(\mathbf{X}_t, \mathbf{p}_t) - \mathbf{X}_{tj} \frac{\partial g_i}{\partial X_j}(\mathbf{X}_t, \mathbf{p}_t) - \mathbf{p}_{tk} \frac{\partial g_i}{\partial \mathbf{p}_{occ_k}}(\mathbf{X}_t, \mathbf{p}_t) \right) + \mathbf{X}_j \frac{\partial g_i}{\partial X_j}(\mathbf{X}_t, \mathbf{p}_t) + \mathbf{p}_{occ_k} \frac{\partial g_i}{\partial \mathbf{p}_{occ_k}}(\mathbf{X}_t, \mathbf{p}_t), \quad (6.20)$$

where we compute the linearisation at the point $(\mathbf{X}_t, \mathbf{p}_t)$. \mathbf{p}_t has two components for the two scalae in this case, and k is an index over these two components. Thus $\mathbf{k}_{1sv} = \frac{\partial g_i}{\partial \mathbf{p}_{occ_k}}(\mathbf{X}_t, \mathbf{p}_t)$, and the balance of the terms will contribute to k_{2sv} . The discussion in section 5.2.3 concerning the dependence of the linearisation on the pressure and state applies to the output linearisation as well. This linearisation must be provided to the fluid model.

Once again, the `update_matrices` method can be used to update the linearisation. To avoid unnecessary calls to this method where the output equation is linear, a matrix equation is implemented to compute \mathbf{k}_1 and k_2 , for example:

$$a_{sv} = (\mathbf{C}_{sv} \mathbf{X}_{dynamic} + D_{sv}) + (\mathbf{E}_{sv} \mathbf{X}_{dynamic} + \mathbf{F}_{sv}) \mathbf{p}_{occ}, \quad (6.21)$$

$$= k_{2sv} + \mathbf{k}_{1sv} \mathbf{p}_{occ}, \quad (6.22)$$

where \mathbf{C}_{sv} , D_{sv} , \mathbf{E}_{sv} , and \mathbf{F}_{sv} are matrices that can be updated as required by the call to `update_matrices`. Of these matrices, \mathbf{C}_{sv} and \mathbf{F}_{sv} are required in most models, and D_{sv} and \mathbf{E}_{sv} are not needed as often. The terms \mathbf{k}_{1sv} and k_{2sv} are accessed through separate function calls because they are required by the the fluid model at different times.

6.3.2 Other Inputs to Output Linearisation

6.3.2.1 External Inputs

When external inputs need to be supplied to the cochlear model, the same mechanisms can be used to incorporate their effects. If they affect the dynamic model, they can either be included into the `dynamic_states_gradient_estimate` method directly, or they can be used to compute a new linearisation using the `update_matrices` method. Once again, the Taylor expansion can be used to generate the linearisation, although other methods of linearising the model are unlikely to be problematic, because the feedback path through the external inputs is not critical to the linear solver.

6.3.2.2 Inputs from Other Segments

Section 4.4.3 discussed the need for inputs from one `OrganOfCortiComplexSegment` object to another. The feedback might arise from the state or the fluid input \mathbf{p}_{occ} of the first segment. It might

affect the state gradient or the output linearisation of the second segment. There are three important cases to consider:

- The state of segment one affects the state gradient or output linearisation of segment two;
- \mathbf{p}_{occ} of segment one only affects the state gradient of segment two;
- \mathbf{p}_{occ} of segment one affects the output linearisation of segment two.

The first two cases can be handled conveniently using the mechanisms discussed already, because the output linearisation at a given time instant is unaffected by the input. (The state at a time point is unaffected by the fluid solver, so the output linearisation also remains unchanged, and the state gradient does not cause an instantaneous change in the state either.) In case three, the output linearisation must be augmented with additional pressure inputs to make the effect of the input on the output linearisation explicit:

$$a_{sv} = k_{2,sv} + \mathbf{k}_{1sv}\mathbf{p}_{occ} + \sum_{k=1}^{n_{coupled\ segments}} \mathbf{k}_{sv_k}\mathbf{p}_{occ_k}, \quad (6.23)$$

where \mathbf{k}_{sv_k} has a structure similar to \mathbf{k}_{1sv} , with constants appropriate for the coupling between the current segment and the k^{th} input segment, and \mathbf{p}_{occ_k} is \mathbf{p}_{occ} for the k^{th} input segment.

Implementing mechanisms to facilitate these types of models requires a few things. Firstly, data structures need to be provided in the `OrganOfCortiComplexSegment` class to describe the functions required. Secondly, an efficient method of accessing state and input to each segment in a way that does not violate encapsulation of the data must be provided. More than one segment might require input from a specific segment. (Currently a single copy of the data is taken by the `OrganOfCortiComplexCompilation` class, and all the segments access this data via pointers as required.) These features have already been implemented in the framework.

A method to specify where each segment should receive inputs from is still required. It is envisioned that the user could supply an object specifying instructions on how to “link” the segments. The instruction could be as simple as “this segment must receive input from the basal segment” or a more complicated rule based on the location of the segment, or even probabilistic considerations. Furthermore, methods to incorporate the \mathbf{k}_{sv_k} values into the \mathbf{K}_{aug} matrix must be provided. This required the use of global degree of freedom indices because \mathbf{p}_{occ_k} may not lie on the same cell the segment (also see section 6.2.3.1). Neither of these appear to be especially complicated to implement, but they have not been required for models implemented in the framework to date. To avoid defining a specification that ultimately proves to be unhelpful, we are waiting until a model requires them before proceeding in this direction.

6.3.3 Data Outputs

In the above sections, the output linearisation and the dynamic states update equation have been discussed. The dynamic states are chosen as required by the fluid-structure feedback loops. In order to obtain information about the state of the organ of Corti, and to monitor the predicted behaviour of parts of the model that are not required for the fluid-structure interaction, data outputs are provided. Some of these outputs may be used as inputs to loops controlling some of the external inputs described above, and others may be used to verify the model. The data states update equations, in their general nonlinear form, are given below:

$$\dot{\mathbf{X}}_{\text{data}} = [f_i(\mathbf{X}_{\text{dynamic}}, \mathbf{X}_{\text{data}}, \mathbf{p}_{occ}, t)]. \quad (6.24)$$

The data outputs for the segment, \mathbf{Y} , in their general nonlinear form are given by:

$$\mathbf{Y} = [g_i(\mathbf{X}_{\text{dynamic}}, \mathbf{X}_{\text{data}}, \mathbf{p}_{occ}, t)]. \quad (6.25)$$

In the base class for organ of Corti segments, these equations are implemented in linear form only. However, as with the output linearisation, the data state update method is marked `virtual`, and the `update_matrices` method may also be used to introduce a nonlinear model.

A single state vector for the model could be made by augmenting the dynamic states to the data states. The dynamic model and the data model would then also be augmented to give a single update and output equation. This is inefficient, if provision is made for a linear state space model to represent the system, because the matrix relating state gradient to the state is a square matrix of dimension n_{states} . Assume that there are $n_{dynamic}$ dynamic states and n_{data} data states. Then $(n_{dynamic} + n_{data})^2$ values are required to store the augmented matrix. On the other hand $n_{dynamic}^2 + n_{data}(n_{dynamic} + n_{data})$ values are required to store the matrices for the dynamic states and data states separately. Comparing these number we see:

$$\begin{aligned} (n_{dynamic} + n_{data})^2 &= n_{dynamic}^2 + 2n_{dynamic}n_{data} + n_{data}^2 \\ &> n_{dynamic}^2 + n_{dynamic}n_{data} + n_{data}^2 \\ &= n_{dynamic}^2 + n_{data}(n_{dynamic} + n_{data}). \end{aligned} \quad (6.26)$$

This shows that separating out the dynamic states and data states reduces the storage required for the segment model. (Similarly, it is better to keep the external inputs and input from other segments separate, rather than using them to augment the dynamic model.) Although the saving per segment might be small, there are a large number of segments in a cochlear model, and it will add up.

6.3.4 Inner Hair Cell and Nerve Models

Section 4.4.4 has discussed the interface between the cochlea and the afferent nerve pathways in some detail. Justification for including the inner hair cells in the cochlear model but excluding the afferent nerves is given in that section. Output from the hair cell model to the neural pathways can make use of the data outputs mechanism described above.

6.3.5 Locating Required Inputs and Outputs

So far we have assumed that the `OrganOfCortiComplexCompilation` class is able to track all segments when inputs are provided for them and outputs requested from them. Early versions of the class used the node numbers assigned by `deal.II` internally to keep track of segments. In the current version of the class, the location of the segment along the basilar membrane boundary with the scala tympani is used instead. The `OrganOfCortiComplexCompilation` uses `deal.II` methods to extract the spatial location of the required point (this is obtained directly from `deal.II` objects that are easily available in the fluid model) and then converts this to a unique `double` using a mapping that the user provides. Including the mapping means that the power of the finite element method to handle complex geometries is not lost. The resulting `double` is used as the key to a `std::map` that provides the index of the correct segment in a `std::vector` of `OrganOfCortiComplexSegments`. Although the location is a `double`, it can be reliably used as the key for the `std::map` class because the mesh remains unchanged as the simulation progresses. The same values can be used to identify the correct segment when incorporating external inputs or providing feedback to the user.

6.3.6 Conventions Regarding Data Outputs and External Inputs

Having discussed the intended use of data outputs and external inputs in the organ of Corti model, the interface and suggested conventions for the model are laid out here. Data outputs and external inputs to individual organ of Corti segments are specified to be of type C++ `std::vector<double>`, but the user will generally access these variables through an organ of Corti compilation.

The `OrganOfCortiComplexCompilation` class provides a method `get_data` which returns a 2D array of data type `std::vector<std::vector<double>>` called `compiled_data`, and a `std::vector<double>` named `keys`. `keys` is the list of segment locations along the basilar membrane boundary with the scala tympani, described above, and `compiled_data` is the data from the segments, as defined by the segment class used. The order of segments is the same in `keys` and `compiled_data`. Therefore, for a given index, `i`, `keys[i]` is the segment location, and `compiled_data[i]` is a `std::vector<double>` of the segment data.

The method `set_external_inputs` in the `OrganOfCortiComplexCompilation` class takes two inputs, a `std::vector<double>` list of segment locations that the user wishes to change the inputs

for (similar to `keys` above), and a `std::vector<std::vector<double>>` of new inputs. The rows of the inputs must be in the same order as the segment locations list.

It is anticipated that the user will first call the `get_data` method, and pass the output to a method which uses the locations and data to derive the next inputs to be applied for a given location. Having computed the inputs, the method would then list the locations that need to be updated and the data required for each location, producing the input required by the `set_external_inputs` method.

6.4 Other Boundary Conditions

Sections 4.5 and 4.6 in the specification discussed how the various boundary conditions may be handled at length. In this chapter, section 6.2.3.2 provided details of how these boundary conditions are incorporated into the fluid model. The state space models required to implement other boundary conditions are simpler than those required for the organ of Corti complex boundary conditions, which have been described in detail in section 6.3. Therefore, at this stage we simply direct the reader to these sections.

6.5 Summary

This chapter has unpacked some of the details required to implement the specification provided in chapter 4. It takes into account the finding of the studies related to the fluid-structure interaction described in chapter 5. Following the pattern of chapter 4, we began by considering the fluid model and then moved on to the structures, in particular the cochlear partition. For the next two chapters we shift from specifying the structure of the proposed framework to providing examples of its use. The applications have been selected to demonstrate the capabilities of the framework, while investigating subjects of interest to the field as a whole.

Chapter 7

A First Application

We have outlined the specification of a cochlear modelling framework that can simulate a wide range of models, and detailed how it can be implemented. We now move on to an example of the use of the framework. A simple passive (linear) model is considered for clarity, but it is obviously preferable if the results from the model are relevant to a current question in cochlear modelling. Two recent publications ([Bell 2010](#) and [Babbs 2011](#)) provide an ideal opportunity. They both develop a common argument against the travelling wave hypothesis: that the radial segments in the cochlea act like isolated resonators, receiving input in parallel, and therefore any observed travelling wave might just be an epiphenomenon. An investigation into these papers suits our purposes, because it can be addressed with a linear model, and it will contribute to an ongoing discussion.

This chapter aims to provide a balanced assessment of the above argument against the travelling wave hypothesis. To provide context, the introduction considers the “failure modes” that any cochlear model is susceptible to, and then reviews some of the other criticisms raised against the travelling wave hypothesis. Section 7.2 develops a mathematical result that completes our study of the fluid-structure interaction in the cochlea, and allows us to comment on the criticism. Section 7.3 describes how the model was implemented in the framework, and section 7.4 presents numerical results from the model. The numerical results confirm the mathematical derivation, and we discuss their implications for both the publications cited ([Bell 2010](#) and [Babbs 2011](#)), and the framework.

7.1 Introduction

7.1.1 Models with Emergent Travelling Wave Behaviour

Some abstract models may explicitly specify a travelling wave, or implicitly exclude other modes of behaviour. Many do not presuppose one form of behaviour or another, and any travelling wave that is visible in the results is emergent behaviour (a desirable characteristic, according to this thesis).

Models that show emergent travelling waves should not really be referred to as “travelling wave models”, because this would imply they exclude other modes of behaviour. We will avoid the terms “travelling wave models” and “resonance models”, preferring to refer instead to the “travelling wave hypothesis” and the “resonance hypothesis”.

The travelling wave has become a conceptual tool for researchers to understand how various phenomena in the cochlear response arise. While these “mental models¹” were established based on physiological and numerical results (no one anticipated Von Békésy’s observation of the travelling wave), they now do presuppose a travelling wave. Certain criticisms that show inadequacies in our “mental models” of the travelling wave have little bearing on models with travelling waves as emergent behaviour.

7.1.2 Ways Technology Projects Can Fail

In science and engineering, widely accepted theories can be disproved and promising technologies can fail. Although physics is able to set hard limits on the possible performance of some systems (for example, the efficiency that a combustion engine can attain), researchers often have to rely on their own intuition about whether a new idea is likely to succeed or not. We now look at how technology projects may fail, because the field of cochlear modelling is essentially a technology project, and given the diversity of conflicting ideas, some of these ideas are guaranteed to fail. Projects that will ultimately fail can fit into two broad categories that will be discussed below. Approaches to cochlear modelling may also fail for a completely different reason, unrelated to technology projects in general: new research may invalidate an otherwise viable model.

7.1.2.1 New Research Invalidates a Key Assumption or Prediction

Cochlear models encapsulate theories of hearing, so they can be invalidated in the same way as any other scientific theory. Theories use assumptions and make testable predictions. If new results either invalidate the assumptions, or are in conflict with the predictions, then the theory is invalidated. The theory must be revised to account for the new data or be abandoned. However, we must be careful not to discard ideas too quickly. We have seen that reliable data from the cochlea is hard to obtain. Furthermore, in something as complex as the cochlea it is hard to control all the possible interactions between the elements. (This is not a criticism of the state of experimental research, but an observation that we need to understand the constraints under which experimental research is

¹Senge (1990) describes mental models: in a response to our limited ability to handle large volumes of data, humans form mental models to generalise it to simpler trends. If we are aware of this process, we can deliberately check the consistency of our mental models against data in order to improve them. This process is clearly as important in science as in the business environment, for which Peter Senge writes.

conducted.) Therefore, when we find data that seems to invalidate a model, we must also be willing to consider that we may not fully understand the data.

Currently, there is no widely accepted model that provides a comprehensive understanding of the cochlea, because *all* models have either been tested to the extent that disagreement with experimental data at some point becomes apparent, or else they have not been tested sufficiently to reach this point. Disagreement with the data at one point presents an opportunity to revise an otherwise promising model, but revision required might require a chance insight by a future researcher. Therefore, the model should not be discarded too quickly, and evaluating models inevitably requires experience and judgement on the part of the researcher.

Another important use of cochlear models is to inspire technology. While indisputable evidence that invalidates a cochlear model would be fatal to that model for research into hearing, it may not affect its use in technology. If the “invalidated” model continues to perform well in important applications it can (and should) continue to be developed and used for these purposes. Although the physiological basis for the model may now be tenuous, if the results justify continued use in a technology application, then there is little point in changing.

7.1.2.2 Iterative Improvements Become Increasingly Prohibitive

Assume we have a candidate for the “ideal solution” of chapter 2. This cochlear model provides a reasonable match to some fraction of the set of data available from the biological cochlea. The candidate solution does not match all the data, but this might be just because researchers have not tested it against that data yet. With further work and refinements, we are able to increase the fraction of data that is matched and the quality of the match overall.

This picture could be characteristic of progress towards the ideal solution. Unfortunately, it might also be wasted effort that will ultimately lead to failure. A model that allows for iterative improvements is ideal for research, but it does not guarantee that the approach is actually sound. If it is not sound, then the amount of improvement made will peter out gradually with increasing effort. The critical question is how much effort is expended to achieve a given amount of return. Past successes and the ongoing opportunities for work can obscure signs that new work is yielding little improvement. By the time the project actually fails, people have expended more resources than are really justified on the idea. In fact, the idea may never be finally disproved, but may simply be abandoned instead. However, many good ideas were abandoned the first time they were tried, but ultimately proved successful at a later stage. Once again, the experience and judgement on the part of the researcher is required.

7.1.2.3 Idea is Untestable Without a Complex, but Key, Set of Components

Some legitimate ideas do not lend themselves to iterative development, because they require a key set of components that do not work in isolation. An example of this from biology is the relationship between RNA (ribonucleic acid) and DNA (deoxyribonucleic acid). DNA provides information storage for a cell and RNA is a messenger to tap into the data. Without RNA, the information in the DNA is inaccessible, and without DNA, the RNA has nothing to carry. Systems with this sort of dependency require all the critical elements to be in place before any testing can be done.

If at least one of the necessary components is difficult to implement, we have a new “failure mode”. These models require a lengthy development period before realistic tests can be performed. In the interim, the performance of the system might be estimated, but these estimates are not necessarily accurate. If the model is still a failure at the end of the lengthy development period, then there are three possible scenarios: the approach may actually be invalid; some aspect of the model may require an even more complex representation; or it may all be sound, with a “bug” somewhere in the complicated implementation. Therefore, even failure tends to drive further work in this “mode”.

There are obviously cases where the resources expended on putting together a set of key components that cannot work in isolation leads to a major breakthrough in performance, so we cannot immediately discount an idea based on its complexity², and a researcher’s insight into the problem is required.

7.1.2.4 Cochlear Modelling and Failure Modes

Every cochlear model can be invalidated by research that disproves either its assumptions or predictions, although we have noted that models do need some latitude to develop, so this test should be applied with care. The predominant view (mental model) on cochlear modelling currently is the travelling wave hypothesis. This is the classical view that has been presented (implicitly to some extent) in chapter 3. Given the regular publications, and the wide variety of variations of models that use this general approach, it is certainly undergoing iterative improvement. If the travelling wave hypothesis ultimately proves to be a failure, then it would have certainly failed in the “iterative” manner, although there is the possibility that some individual developments based on the travelling wave hypothesis may fail in the “key component” manner.

We have noted that some researchers dispute the travelling wave hypothesis. To the extent that the alternatives these researchers put forward can be tested iteratively, they can succeed or fail in that mode. Some of these researchers have expressed to this author the view that unless certain key parts (that are fairly complex) are correctly represented, their model would be incomplete and would not

²Just because an idea that is complex for us to implement, does not mean it is complex for the cochlea to use. Section 4.2 discussed the difficulties associated with numerical modelling of the cochlear fluid properties – a problem the biological cochlea does not face.

work³. To the extent that this is correct, these models will succeed or fail in the “key component” mode. Tests that estimate the performance of these models help to reduce the risks associated with this failure mode, but they do not change it into the “iterative” mode.

7.1.3 Common Criticisms of the Travelling Wave Hypothesis

Certain criticisms are often raised against the travelling wave hypothesis. Understandably, these criticisms are often raised by people who suggest that the travelling wave hypothesis should be abandoned. They are also sometimes pointed out by researchers who are strong advocates of the travelling wave hypothesis. “Self-criticism” within the modelling community is a sign of a healthy research field that is able to identify research problems that remain interesting and controversial, and should be treated as such.

Andrew Bell listed 13 “anomalies in travelling wave theory” in his PhD thesis (Bell 2005, page I3[21]). Drawing on this reference amongst others, a list of five criticisms that are repeatedly raised in mailing list discussions is presented. The aim of this list is to acknowledge that there are areas where current travelling wave theory is incomplete – providing opportunities for further research into travelling wave theory or alternatives, depending on the researcher’s intuition about which approach is likely to be more fruitful. While these areas are still unexplained, it seems reasonable to also consider the case for alternative models that can address specific areas, and a platform that allows uncommon ideas of how the cochlea works to be tested alongside the more popular ideas will be useful in this process.

7.1.3.1 Size of Signals that must Stimulate the Basilar Membrane

The ear is remarkably sensitive. Braun (1994) states that the energy input required to trigger sensation in the ear is just 5×10^{-11} erg⁴ at 1 kHz, which is even more sensitive than the eye (the eye requires 1.3 to 2.6×10^{-10} erg). We are able to hear a sound that causes a stapes displacement of just a tenth of a diameter of a hydrogen atom. The magnitude of the input displacement and energy that can cause a hearing sensation is something that all models strive to account for. It is also a reason that the ear is of interest to engineers who aim to emulate its capabilities.

In travelling wave theory, the travelling wave “funnels” energy to the point that is most sensitive to its particular frequency, possibly enhancing the signal along the way using the cochlear amplifier –

³This author is sympathetic to such statements: many researchers study the travelling wave hypothesis, but only a few study the alternatives. A single negative report in the literature based on halfhearted or incomplete work by someone who is not committed to the model may be very damaging to an upcoming model. This is part of the reason that I *will not* come to any conclusion about the validity of any specific models in this work. I am concerned with creating a framework that is fair to a wide variety of models, and will leave the actual comparisons to future researchers who focus on a smaller subset.

⁴1 erg = 1 g · cm² · s⁻² = 10⁻⁷ J.

seemingly an efficient way to use the input energy. The criticism of this scenario is that the size of the basilar membrane vibration is also small. Bell states that experimental evidence shows that the displacement of the basilar membrane is only 1 nm at the threshold of hearing (Bell 2010). He goes on to estimate that if the cochlear amplifier provides a gain of 10^3 , the cochlea must be detecting a displacement on the order of picometres in a “travelling wave” model. This calculation assumes that the cochlear amplifier only begins to function once a passive signal has been detected. Chapter 8 addresses a class of models with emergent travelling wave behaviour that require active mechanisms to be constantly tuning the cochlea to the point of oscillation (a Hopf-bifurcation) via local and neural feedback loops. The Hopf-bifurcation type of model avoids the necessity to detect signals on the order of picometres.

A related area of criticism for active cochlear models concerns the high gains that they sometimes entail. Neely and Kim (1986) state values as high as $30000\times$ for the “total power gain” (which they define as the ratio of the power absorbed by the cochlear partition, over all regions where the driving-point-impedance has a positive real part, to the power input at the stapes). Bell criticises these high gains because “the actual cochlea must make good use of all the signal energy available” (Bell 2005, page I3 [24]). These gains would only be required for the smallest input signals, so the actual energy use implied is only 1.5×10^{-6} erg or about 3.6×10^{-17} kcal⁵. The correct question is not whether the ear has used the input energy as efficiently as possible, but whether the energy expenditure is worth it for the ability to hear sounds that lie at, or even below, 0 dB (SPL).

Andrew Bell and Martin Braun are two authors who suggest that the outer hair cells are the first recipients of signals entering the ear (Bell 2005; Braun 1994). They propose that the outer hair cells are able to detect the fast pressure wave in the cochlea and amplify the signal for the inner hair cells. They think that it is more plausible for the small outer hair cells to respond to the weak input pressure than for a vibration pattern to be setup on the comparatively large basilar membrane. However, at the time of writing, it has not been conclusively proven that outer hair cells are sensitive to pressure directly.

7.1.3.2 Range of Frequencies that must be Covered by the Cochlea

The undamped resonance frequency of the passive basilar membrane model, given in equation (4.33), is determined by:

$$f(x) = \frac{1}{2\pi} \sqrt{\frac{K(x)}{M(x)}}. \quad (7.1)$$

The square root that appears in the equation is inconvenient in the sense that it means that for $f(x)$ to vary by three orders of magnitude (20 Hz to 20 kHz), the ratio $K(x)/M(x)$ must vary by six. Bell (2005), section 3.2/b, reviews some literature showing that there has been ongoing debate over

⁵1 kcal = 4.184×10^3 J. For comparison, healthy daily calorie intakes are of the order of 10^3 kcal.

whether $K(x)/M(x)$ can vary over a sufficient range. He concludes that “real doubts remain about being able to achieve a satisfactory range of tuning”.

The first cochlear modellers were aware of the difficulties that would be encountered when accounting for this large variation required but, on the balance of evidence available, thought these models were viable (see [Lighthill 1981](#) for a review). Although the stiffness of the basilar membrane needs to vary by a large amount, this does not necessarily imply that its physical dimensions must change by the same amount: the stiffness of the basilar membrane is proportional to the third power of its height, and the fifth power of its width ([Cole and Chadwick 1977](#)). Recently Charles Babbs also assessed the dynamics of basilar membrane, and found reasonable agreement between his predicted cochlear map and that in [Greenwood \(1990\)](#) (see [Babbs 2011](#)). (We will dispute his assumption that radial segments act in isolation, but this does not affect the stiffness values he estimates.)

In this chapter, we will show that equation (7.1) is too simplistic to use to tune the cochlear partition directly if the mass of the partition $M(x)$ is low. Nevertheless, the stiffness $K(x)$ must be able to vary over a wide range for the travelling wave hypothesis to be plausible. While Bell raises some points that interested readers can follow up, the theoretical results discussed suggest the necessary range can be attained.

7.1.3.3 Sharpness of the Cochlear Response

The neural response of the ear to sound is remarkably sharp. Whether the response of the basilar membrane is as sharp is open to debate. Figure 3.13 showed a correlation between the basilar membrane tuning curve and the neural turning curve that would suggest that it is. However, Jont Allen notes that this is a key question for the field: “The discrepancy in frequency selectivity between basilar membrane and neural responses has always been, and still is, the most serious problem for the cochlear modelling community. *In my view, this discrepancy is one of the most basic unsolved problems of cochlear modelling.*” ([Allen 2001](#), italics appeared in the original text.) He continues that uncertainty in, and the interpretation of, the experimental data has hindered progress in addressing this issue.

There are two separate questions: (a) should the basilar membrane tuning be as sharp as the neural tuning; and (b) can the travelling wave hypothesis provide basilar membrane tuning that matches physiological data. The answer to (a) relies primarily on the physiological data. Cochlear models must conform to reliable physiological data. If cochlear models do not need to predict basilar membrane cut off slopes equal to the neural data, this removes some rather strict requirements that would otherwise be placed on the model. However, it does open the question of what causes the additional sharpness in the neural response.

Addressing (b) also requires accurate physiological data, but cochlear models are valuable. They can be used in two separate ways: we can place limits on the structure of the model (the size of

parameters for example) based on our physiological data, and observe whether the response we obtain agrees with the cochlear response; or we can tune our model parameters to obtain a close match to the cochlear response, and observe the implications of the parameters that we required. Solving the “inverse problem” (for example, see [de Boer 1996](#), section 6) where parameters are forced to values that produce the desired results, rather than adjusted by hand to give a fair fit, is an extreme example of the second approach. Criticisms against travelling wave theory concerning the unrealistic values that this process can produce are not really justified – it is simply another way that modellers can work with experimentalists to understand the mechanisms that are active in the cochlea.

Given the flexibility afforded to variations of the travelling wave hypothesis to approach the physiological data from a wide range of directions, it seems only fair to allow alternative models the same breadth of approaches. To this author, it seems that the case for the travelling wave has been strengthened by use of numerical models that allow emergent behaviour, and by frequent critical assessment of the idea by its advocates as well as its critics. I hope that alternative models also get these opportunities!

Finally, using a scanning laser interferometer, [Ren \(2002\)](#) found evidence for a travelling wave on the basilar membrane, but over a restricted area ($\approx 600\mu\text{m}$ for a low level 16 kHz tone). This might seem to contradict travelling wave theory, which predicts a wave on the basilar membrane between the stapes and the best place. Ren found that the range over which the wave could be detected increased with the size of the input, and that it increased more quickly basally than apically, which is consistent with travelling wave theory. Interestingly, the Van der Pol oscillator model, studied in chapter 8, can show similar behaviour to that observed by Ren (see section 8.4 and [Duifhuis *et al.* 1986](#), figures 5 and 6).

7.1.3.4 Simultaneous Delivery of the Input via a Fast Pressure Wave

While some criticisms of the travelling wave hypothesis are based on physiological data, others are based on conceptual arguments. (The previous section is a combination of a conceptual argument and physiological data.) We have noted that the speed of sound in the cochlear fluid is approximately $1500\text{ m}\cdot\text{s}^{-1}$, and it is common to ignore the slight compressibility of the fluid that limits it, meaning that the speed of sound is effectively infinite in most cochlear models. This seems to imply that every point in the cochlea experiences any input at the stapes simultaneously, that is, in parallel. If all parts of the cochlea receive this signal in parallel, why should only the basal region of the basilar membrane respond to it? If there is no structural coupling between the radial segment, where are the forces between neighbouring segments that a travelling wave requires? Perhaps the entire cochlear partition responds to it, and the “travelling wave” we observe in steady state results is merely an epiphenomenon.

There is both experimental evidence and numerical evidence for the existence of a travelling wave on the basilar membrane. Passive cochlear models consisting of a resonant cochlear partition, a coupling fluid, and a differential pressure have been solved in the frequency domain using numerous techniques; for example, the 2D finite difference method (Neely 1981a) and the Liouville-Green approach (Steele and Taber 1979a). If the phase response of these models is “unwrapped” to avoid discontinuities around 2π , a smooth curve is obtained. This curve shows a phase delay that increases with distance from the stapes to figures greatly in excess of 2π . Phase delays of greater than 2π were observed by Von Békésy (see figure 3.10) and were a key indicator that travelling waves were dominating the observed behaviour (as commented by Ren 2002).

While the passive travelling wave predicted by frequency domain cochlear models shows agreement with the phase delay in physiological data, the fact that all parts of the cochlea seem to receive the input in parallel has led to the argument that the travelling wave and its delay may just be an epiphenomenon. This argument has been worked out in some detail in the recent publications by Andrew Bell (Bell 2010) and Charles Babbs (Babbs 2011), previously mentioned. It seems that time domain cochlear models offer a unique method to investigate this question further, because we are able to observe how the steady state travelling wave pattern is established in great detail. Therefore, this criticism of the travelling wave will be the subject of the numerical experiments presented in this chapter.

7.1.3.5 Cochlear Response to Various Pathological Conditions

Section 4.6.3 discussed alternative fluid boundary conditions that can become important when the cochlea is in pathological conditions. The response of the cochlea in these conditions can also give insight into the importance of travelling waves in the cochlea. Travelling waves in a cochlear model require certain features to exist: fluid coupling between segments; partition dynamics that allow resonance⁶; and a differential pressure across the partition as input. Arguments against the travelling wave based on pathological conditions aim to setup a scenario where some of the features required to support travelling waves are missing, yet the sensitivity of the ear is unaffected. This section is “catch-all” for various arguments along those lines.

A number of pathological conditions focus on situations when the differential pressure is presumed to be zero, but the compression wave is still present. Examples include lesions in the wall of the cochlea, a rigid round window condition, and cases where the middle ear is missing. Bell (2005) discusses some examples in further detail. The results are difficult to interpret, because it is hard to say exactly what changes have occurred in the cochlea, especially when the examples are from medical conditions. Tonndorf and Tabor (1962) investigated the closure of the cochlear windows

⁶Strictly, resonance is not required for a travelling wave, but Lighthill (1981) argues that it is required for the “critical-layer absorption” that is seen in the cochlea.

and showed that other signal pathways become important in these cases. Small differences in the conditions can also have surprising results. For example, they discussed how air bubbles accidentally trapped by the round window or in the scala tympani could result in an *improved* response where a degradation in the response was expected. Cases from medical science cannot be ignored – especially because they provide human data for scenarios that would be unethical to obtain via direct experiment – but they remain difficult to interpret. Therefore, these types of conditions require careful consideration and investigation when used as arguments for or against the importance of travelling waves in the cochlea.

A second target for these types of arguments could be the continuity of the cochlear partition. If a discontinuity in the cochlear partition allows the differential pressure to equalise at that point (say a physical hole) this reverts back to the previous argument. If, instead, the differential pressure is maintained but the dynamics of the partition in the region are affected, then this might seem to exclude a travelling wave. (See [Dallos 1996](#), figure 1.8 for a summary of the effects of outer hair cell lesions on nerve fibre responses.) If our “mental model” of the cochlea dictates that all the energy that the passive travelling wave can possess must enter the wave at the basal end of the basilar membrane, and be transferred along the basilar membrane to the best place where it is detected, then these results are certainly a blow to that view. However, we have stressed that the travelling wave (whether it is essential or not) is emergent behaviour due to the interaction of the fluid and cochlear partition. A “dead” region in the cochlea need not stop a travelling wave altogether, because the fluid still couples radial segments on either side of the region. This second criticism is not necessarily a criticism of the function of travelling waves in the cochlea, but of the mental models in the minds of either its proponents, or its critics.

7.2 Segment Coupling through the Cochlear Fluid

The case for resonance that was set out in [Bell \(2010\)](#), and separately in [Babbs \(2011\)](#), revolves around the argument that a set of independently excited resonators can exhibit behaviour that looks like a travelling wave. This argument is essentially the criticism addressed in section 7.1.3.4. We will ultimately investigate this criticism using a numerical model, but we start with an investigation into the nature of the segment coupling through the fluid using mathematical manipulations of the monolithic state space cochlear models presented in section 5.1.1. For simplicity, we assume a single scala cochlear model to derive the required result.

7.2.1 Classical Cochlear Partition Models

Of the two expressions for the monolithic state space model, it is most convenient to start with the approach introduced by [Elliott *et al.* \(2007\)](#) on this occasion. The state update and output equations

are then:

$$\dot{\mathbf{X}}(t) = [\mathbf{I} - \mathbf{B}_E \mathbf{F}^{-1} \mathbf{C}_E]^{-1} \mathbf{A}_E \mathbf{X}(t) + [\mathbf{I} - \mathbf{B}_E \mathbf{F}^{-1} \mathbf{C}_E]^{-1} \mathbf{B}_E \mathbf{F}^{-1} \mathbf{q}(t), \quad (7.2)$$

$$\mathbf{a}(t) = \mathbf{C}_E \dot{\mathbf{X}}(t). \quad (7.3)$$

\mathbf{A}_E , \mathbf{B}_E , and \mathbf{C}_E have a block diagonal structure, for example:

$$\mathbf{A}_E = \begin{bmatrix} \mathbf{A}_1 & 0 & \cdots & & \\ 0 & \mathbf{A}_2 & & & \vdots \\ \vdots & & \ddots & & 0 \\ & \cdots & 0 & \mathbf{A}_{N+1} & \end{bmatrix}.$$

Terms in the first block (\mathbf{A}_1 , \mathbf{B}_1 , and \mathbf{C}_1) correspond to the state space model of the stapes boundary condition, and the remaining N blocks are each associated with a specific location ξ_i on the boundary between the cochlear partition and the fluid region, Γ_{cp} . (We assume that ξ_i are ordered such that distance from the stapes increases with i .) Similarly, values associated with the stapes are at the beginning of $\mathbf{a}(t)$, $\mathbf{q}(t)$, and $\mathbf{X}(t)$.

We now specify the state space matrices \mathbf{A}_i , \mathbf{B}_i , and \mathbf{C}_i that are required to complete the model outlined above. As discussed in section 5.2.3, ‘‘classical’’ models of the cochlear partition define $a(\xi_i, t)$ as a (possibly nonlinear) function of $p(\xi_i, t)$ only. These models are constructed by assuming that the pressure difference across the partition, $2p(\xi_i, t)$, together with any active mechanisms present, act as forcing terms on the passive dynamics of partition to cause an acceleration of the mass of the partition. Therefore, the physics of the problem requires:

$$M(\xi_i) a(\xi_i, t) = f(\xi_i, v(\xi_i, t), d(\xi_i, t), \cdots) + 2p(\xi_i, t), \quad (7.4)$$

where $M(\xi_i)$ is the mass of the partition per unit area. $v(\xi_i, t)$ and $d(\xi_i, t)$ are the velocity and displacement of the partition respectively. While these variables are sufficient, f may also depend on additional parameters. This is a statement of Newton’s second law, applied to the partition as a whole.

We will study linear cochlear models in this work, so the radial segments of the cochlear partition may be modelled by a linear state space system:

$$\dot{\mathbf{X}}(\xi_i, t) = \mathbf{A}(\xi_i) \mathbf{X}(\xi_i, t) + \mathbf{B}(\xi_i) p(\xi_i, t), \quad (7.5)$$

$$a(\xi_i, t) = \mathbf{C}(\xi_i) \dot{\mathbf{X}}(\xi_i, t), \quad (7.6)$$

where the state $\mathbf{X}(\xi_i, t)$ and state gradient $\dot{\mathbf{X}}(\xi_i, t)$ are the $n \times 1$ vectors. $\mathbf{A}(\xi_i)$, $\mathbf{B}(\xi_i)$, and $\mathbf{C}(\xi_i)$ are $n \times n$, $n \times 1$, and $1 \times n$ respectively. $\mathbf{X}(t) = \begin{bmatrix} \mathbf{X}_s^T(t) & \mathbf{X}^T(\xi_1, t) & \cdots & \mathbf{X}^T(\xi_2, t) \end{bmatrix}^T$, given the

ordering of \mathbf{A}_E above.

We note that the physics of the cochlea implies that $\mathbf{A}(\xi_i)$, $\mathbf{B}(\xi_i)$, and $\mathbf{C}(\xi_i)$ will generally have a special structure: $\mathbf{B}(\xi_i)$ and $\mathbf{C}(\xi_i)$ can be expressed as vectors with only one nonzero term, in which case only the first row in $\mathbf{A}(\xi_i)$ will generally be dependent on $M(\xi_i)$. Firstly, define $\dot{\mathbf{X}}_1(\xi_i, t) = a(\xi_i, t)$ then $\mathbf{C}(\xi_i) = \begin{bmatrix} 1 & \mathbf{0} \end{bmatrix}$, and $\mathbf{X}_1(\xi_i, t)$ is the velocity of the partition at ξ_i . The first row of equation (7.5) must then satisfy (7.4):

$$M(\xi_i)a(\xi_i, t) = \tilde{\mathbf{A}}_{1j}(\xi_i)\mathbf{X}_j(\xi_i, t) + 2p(t), \quad (7.7)$$

where $\mathbf{A}_{1j}(\xi_i) = \tilde{\mathbf{A}}_{1j}(\xi_i)/M(\xi_i)$ is also a ‘‘per unit area’’ quantity. Only summation over the matrix index j is implied (summation over segments ξ_i is not required in this derivation). From this equation, it is clear that $\mathbf{B}_1(\xi_i) = 2/M(\xi_i)$. Similarly $\mathbf{A}_{1j}(\xi_i) \propto 1/M(\xi_i)$. The remaining states, $\mathbf{X}_{j \in \{2, \dots, n\}}$, must be linearly independent of X_1 in a state space model. Therefore, unless the pressure has some additional effect on the partition that is not captured by Newton’s second law, $\mathbf{B}_{j \in \{2, \dots, n\}} = 0$. Similarly, $\mathbf{A}_{jk, j \in \{2, \dots, n\}}$ are not dependent on $M(\xi_i)$, because $M(\xi_i)$ is the inertia term associated with $a(\xi_i, t)$, and its effect has already been included.

Then equation (7.5) for ξ_i can be expressed as:

$$\begin{bmatrix} \dot{\mathbf{X}}_1 \\ \dot{\mathbf{X}}_2 \end{bmatrix} = \begin{bmatrix} \frac{\tilde{\mathbf{A}}_{11}(\xi_i)}{M(\xi_i)} & \frac{\tilde{\mathbf{A}}_{12}(\xi_i)}{M(\xi_i)} \\ \mathbf{A}_{21}(\xi_i) & \mathbf{A}_{22}(\xi_i) \end{bmatrix} \begin{bmatrix} \mathbf{X}_1 \\ \mathbf{X}_2 \end{bmatrix} + \begin{bmatrix} \frac{2}{M(\xi_i)} \\ \mathbf{0} \end{bmatrix} p(\xi_i, t), \quad (7.8)$$

$$a_{cp}(\xi_i, t) = \begin{bmatrix} 1 & \mathbf{0} \end{bmatrix} \begin{bmatrix} \dot{\mathbf{X}}_1 \\ \dot{\mathbf{X}}_2 \end{bmatrix}, \quad (7.9)$$

where block notation is adopted to separate the matrices into terms dependent on X_1 and the remaining states $\mathbf{X}_2 = \mathbf{X}_{j \in \{2, \dots, n\}}$.

7.2.2 Isolating Segment Mass

In this section, we use the special structure noted in $\mathbf{A}(\xi_i)$, $\mathbf{B}(\xi_i)$, and $\mathbf{C}(\xi_i)$ to collect this equation into terms depending on $M(\xi_i)$, and those independent of it. In order to do this, we use a permutation matrix \mathbf{P} .

Let $\mathbf{X}'(t) = \mathbf{P}\mathbf{X}(t)$ where:

$$\mathbf{X}'(t) = \begin{bmatrix} \mathbf{X}_s^T(t) & \mathbf{X}_1(\xi_1, t) & \cdots & \mathbf{X}_1(\xi_N, t) \\ \mathbf{X}_j^T(\xi_1, t)|_{j \in \{2, \dots, n\}} & \cdots & \mathbf{X}_j^T(\xi_N, t)|_{j \in \{2, \dots, n\}} \end{bmatrix}^T. \quad (7.10)$$

This defines the mapping \mathbf{P} . \mathbf{P}^T serves as an inverse mapping because $\mathbf{P}^T\mathbf{P} = \mathbf{I}$ for permutation

matrices. Then:

$$\dot{\mathbf{X}}'(t) = \mathbf{P}\dot{\mathbf{X}}(t) = \mathbf{P}[\mathbf{I} - \mathbf{B}_E\mathbf{F}^{-1}\mathbf{C}_E]^{-1}\mathbf{A}_E\mathbf{X}(t) + \mathbf{P}[\mathbf{I} - \mathbf{B}_E\mathbf{F}^{-1}\mathbf{C}_E]^{-1}\mathbf{B}_E\mathbf{F}^{-1}\mathbf{q}(t), \quad (7.11)$$

$$\begin{aligned} &= \mathbf{P}[\mathbf{I} - \mathbf{B}_E\mathbf{F}^{-1}\mathbf{C}_E]^{-1}\mathbf{P}^T\mathbf{P}\mathbf{A}_E\mathbf{P}^T\mathbf{X}'(t) \\ &\quad + \mathbf{P}[\mathbf{I} - \mathbf{B}_E\mathbf{F}^{-1}\mathbf{C}_E]^{-1}\mathbf{P}^T\mathbf{P}\mathbf{B}_E\mathbf{F}^{-1}\mathbf{q}(t), \end{aligned} \quad (7.12)$$

$$\begin{aligned} &= [\mathbf{I} - \mathbf{P}\mathbf{B}_E\mathbf{F}^{-1}\mathbf{C}_E\mathbf{P}^T]^{-1}\mathbf{P}\mathbf{A}_E\mathbf{P}^T\mathbf{X}'(t) \\ &\quad + [\mathbf{I} - \mathbf{P}\mathbf{B}_E\mathbf{F}^{-1}\mathbf{C}_E\mathbf{P}^T]^{-1}\mathbf{P}\mathbf{B}_E\mathbf{F}^{-1}\mathbf{q}(t). \end{aligned} \quad (7.13)$$

Under \mathbf{P} and \mathbf{P}^T we obtain:

$$\mathbf{C}_E\mathbf{P}^T = \begin{bmatrix} \mathbf{C}_s & & \\ & \mathbf{I} & \mathbf{0} \end{bmatrix}, \quad (7.14)$$

$$\mathbf{P}\mathbf{B}_E = \begin{bmatrix} \mathbf{B}_s & & \\ & \mathbf{M}^{-1} & \\ & & \mathbf{0} \end{bmatrix}, \quad (7.15)$$

$$\mathbf{P}\mathbf{A}_E\mathbf{P}^T = \begin{bmatrix} \mathbf{A}_1 & & \\ & \mathbf{M}^{-1}\tilde{\mathbf{A}}_{11} & \mathbf{M}^{-1}\tilde{\mathbf{A}}_{12} \\ & \mathbf{A}_{21} & \mathbf{A}_{22} \end{bmatrix}, \quad (7.16)$$

where $\mathbf{M}_{ij} = M(\xi_i)\delta(i, j)$. The dimensions of $\mathbf{a}(t)$, $\mathbf{q}(t)$, and \mathbf{F}^{-1} relate to the fluid domain rather than the state, so they are not mapped.

We now study the term $[\mathbf{I} - \mathbf{P}\mathbf{B}_E\mathbf{F}^{-1}\mathbf{C}_E\mathbf{P}^T]^{-1}$. In order to do this, we separate \mathbf{F}^{-1} into blocks coupling the stapes $\mathbf{G}_{s,s}$, the cochlear partition $\mathbf{G}_{cp,cp}$, and between the two $\mathbf{G}_{s,cp}$, $\mathbf{G}_{cp,s}$. (The change in letter stresses the fact that these matrices are not necessarily invertible. $\mathbf{G}_{s,cp}$, $\mathbf{G}_{cp,s}$ are generally not even square.) We find:

$$\begin{aligned} [\mathbf{I} - \mathbf{P}\mathbf{B}_E\mathbf{F}^{-1}\mathbf{C}_E\mathbf{P}^T]^{-1} &= \left[\mathbf{I} - \begin{bmatrix} \mathbf{B}_s & & \\ & \mathbf{M}^{-1} & \\ & & \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{G}_{s,s} & \mathbf{G}_{cp,s} \\ \mathbf{G}_{s,cp} & \mathbf{G}_{cp,cp} \end{bmatrix} \begin{bmatrix} \mathbf{C}_s & & \\ & \mathbf{I} & \mathbf{0} \end{bmatrix} \right]^{-1} \\ &= \begin{bmatrix} \mathbf{I} - \mathbf{B}_s\mathbf{G}_{s,s}\mathbf{C}_s & -\mathbf{B}_s\mathbf{G}_{cp,s} & \\ -\mathbf{M}^{-1}\mathbf{G}_{s,cp}\mathbf{C}_s & \mathbf{I} - \mathbf{M}^{-1}\mathbf{G}_{cp,cp} & \\ & & \mathbf{I} \end{bmatrix}^{-1}. \end{aligned} \quad (7.17)$$

The blocks $-\mathbf{M}^{-1}\mathbf{G}_{cp,s}\mathbf{C}_s$ and $-\mathbf{B}_s\mathbf{G}_{s,cp}$ are full, so it is clear that the stapes is providing an additional feedback path between segments of the cochlear partition. These additional pathways limit our ability to study the effect of \mathbf{M} on the system, leaving us two options: we can express the stapes state space model in the structure of equations (7.8–7.9), which would be true to the physics be-

cause Newton's second law also governs the stapes motion; or we can remove the stapes model and enforce an acceleration $\mathbf{q}_1(t)$ at the stapes, so that $\mathbf{a}_1(t) = 0$. The option of removing the stapes model is acceptable because we are working with a linear cochlear model. We adopt this approach for simplicity, and remove rows and columns relating to feedback paths that are no longer required.

Equation (7.2) becomes:

$$\begin{aligned} \dot{\mathbf{X}}'(t) = & \begin{bmatrix} \mathbf{I} - \mathbf{M}^{-1}\mathbf{G}_{cp,cp} & \\ & \mathbf{I} \end{bmatrix}^{-1} \begin{bmatrix} \mathbf{M}^{-1}\tilde{\mathbf{A}}_{11} & \mathbf{M}^{-1}\tilde{\mathbf{A}}_{12} \\ \mathbf{A}_{21} & \mathbf{A}_{22} \end{bmatrix} \mathbf{X}'(t) \\ & + \begin{bmatrix} \mathbf{I} - \mathbf{M}^{-1}\mathbf{G}_{cp,cp} & \\ & \mathbf{I} \end{bmatrix}^{-1} \begin{bmatrix} \mathbf{M}^{-1} \\ \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{G}_{cp,s} & \mathbf{G}_{cp,cp} \end{bmatrix} \mathbf{q}(t). \end{aligned} \quad (7.18)$$

Noting that $[\mathbf{I} - \mathbf{M}^{-1}\mathbf{G}_{cp,cp}]^{-1} = [\mathbf{M} - \mathbf{G}_{cp,cp}]^{-1}\mathbf{M}$, we find:

$$\begin{bmatrix} \mathbf{I} - \mathbf{M}^{-1}\mathbf{G}_{cp,cp} & \\ & \mathbf{I} \end{bmatrix}^{-1} = \begin{bmatrix} [\mathbf{M} - \mathbf{G}_{cp,cp}]^{-1}\mathbf{M} & \\ & \mathbf{I} \end{bmatrix}, \quad (7.19)$$

where we use the fact that the matrix is block diagonal by taking the inverse of diagonal blocks individually. Using this in equation (7.18) we obtain our desired relationship:

$$\begin{aligned} \dot{\mathbf{X}}'(t) = & \begin{bmatrix} [\mathbf{M} - \mathbf{G}_{cp,cp}]^{-1}\tilde{\mathbf{A}}_{11} & [\mathbf{M} - \mathbf{G}_{cp,cp}]^{-1}\tilde{\mathbf{A}}_{12} \\ \mathbf{A}_{21} & \mathbf{A}_{22} \end{bmatrix} \mathbf{X}'(t) \\ & + \begin{bmatrix} [\mathbf{M} - \mathbf{G}_{cp,cp}]^{-1} \\ \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{G}_{cp,s} & \mathbf{G}_{cp,cp} \end{bmatrix} \mathbf{q}(t). \end{aligned} \quad (7.20)$$

In this expression, the dependence of the state gradient on the mass of the segments \mathbf{M} is explicitly shown. In the next section, we will consider the effect of increasing and decreasing \mathbf{M} , while leaving other quantities (most importantly the fluid density that contributes to $\mathbf{G}_{cp,cp}$) unchanged.

7.2.3 Behaviour for Mass Limits

In this section we first consider the behaviour of (7.20) for the case when terms in \mathbf{M} are large, and then consider the case when they are small.

7.2.3.1 High Mass Scenario

If $\min_i (M(\xi_i)) \gg \max_{i,j} |(\mathbf{G}_{cp,cp})_{ij}|$, then $[\mathbf{M} - \mathbf{G}_{cp,cp}]^{-1} \approx \mathbf{M}^{-1}$, and from (7.20):

$$\begin{aligned} \dot{\mathbf{X}}'(t) &\approx \begin{bmatrix} \mathbf{M}^{-1}\tilde{\mathbf{A}}_{11} & \mathbf{M}^{-1}\tilde{\mathbf{A}}_{12} \\ \mathbf{A}_{21} & \mathbf{A}_{22} \end{bmatrix} \mathbf{X}'(t) + \begin{bmatrix} \mathbf{M}^{-1} \\ \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{G}_{cp,s} & \mathbf{G}_{cp,cp} \end{bmatrix} \mathbf{q}(t) \\ &= \mathbf{P}\mathbf{A}_E\mathbf{P}^T \mathbf{x}'(t) + \mathbf{P}\mathbf{B}_E \begin{bmatrix} \mathbf{G}_{cp,s} & \mathbf{G}_{cp,cp} \end{bmatrix} \mathbf{q}(t). \end{aligned} \quad (7.21)$$

Converting back to the variables $\dot{\mathbf{X}}(t)$ and $\mathbf{X}(t)$ we obtain:

$$\dot{\mathbf{X}}(t) = \mathbf{A}_E \mathbf{X}(t) + \mathbf{B}_E \begin{bmatrix} \mathbf{G}_{cp,s} & \mathbf{G}_{cp,cp} \end{bmatrix} \mathbf{q}(t). \quad (7.22)$$

This is equivalent to the collected form of equation (7.5), with $\begin{bmatrix} \mathbf{G}_{cp,s} & \mathbf{G}_{cp,cp} \end{bmatrix} \mathbf{q}(t)$ providing the input to each segment, replacing $p(\xi_i, t)$. Only $\mathbf{q}_1(t)$ is nonzero, so $(\mathbf{G}_{cp,s})_i$ defines the gain that the i^{th} segment receives from the fluid. The fluid behaviour is governed by the Laplace equation, and if the fluid domain is rectangular, with the stapes and helicotrema on opposite faces, then values in $\mathbf{G}_{cp,s}$ follow the trend: $(\mathbf{G}_{cp,s})_i \approx (\mathbf{G}_{cp,s})_1 + m(|\xi_1 - \xi_i|)$, where $(\mathbf{G}_{cp,s})_1 < 0$, $m > 0$, and $(\mathbf{G}_{cp,s})_N \rightarrow 0$ for N large.

Equation (7.22) shows that the cochlear partition behaviour approaches that of a set of isolated resonators in the high mass scenario, with only the input gain defined by the fluid region. The magnitude of the input gain decreases monotonically from the stapes to the helicotrema.

7.2.3.2 Low Mass Scenario

If $\min_i |(\mathbf{G}_{cp,cp})_{ii}| \gg \max_i (M(\xi_i))$, then $[\mathbf{M} - \mathbf{G}_{cp,cp}]^{-1} \approx -\mathbf{G}_{cp,cp}^{-1}$. For a 1D fluid region discretised using the finite difference method, $\mathbf{G}_{cp,cp}^{-1}$ can be formed explicitly even in the limit $M \rightarrow 0$. $\mathbf{G}_{cp,cp}^{-1}$ arises from the Laplace operator, d^2/dx^2 in 1D. Therefore, it is a tridiagonal matrix. When the finite element method is used to discretise the fluid, the situation is slightly more complicated, because the regularisation provided by \mathbf{M} is necessary for $[\mathbf{M} - \mathbf{G}_{cp,cp}]^{-1}$ to be full rank and invertible⁷. However, $\mathbf{G}_{cp,cp}^{-1}$ still arises from the Laplace operator in the finite element case. Assuming the 1D

⁷The finite element replacement for \mathbf{F}^{-1} proposed in Rapson *et al.* (2012) defines segments at quadrature points on Γ_{cp} , and interpolates values at nodes on Γ_{cp} to these quadrature points. This interpolation matrix is rank deficient, making the replacement as a whole rank deficient, and therefore not invertible in isolation.

finite difference case so that we can take the limit, we obtain:

$$\dot{\mathbf{X}}'(t) \approx \begin{bmatrix} -\mathbf{G}_{cp,cp}^{-1}\tilde{\mathbf{A}}_{11} & -\mathbf{G}_{cp,cp}^{-1}\tilde{\mathbf{A}}_{12} \\ \mathbf{A}_{21} & \mathbf{A}_{22} \end{bmatrix} \mathbf{X}'(t) \quad (7.23)$$

$$+ \begin{bmatrix} -\mathbf{G}_{cp,cp}^{-1} \\ \mathbf{0} \end{bmatrix} \begin{bmatrix} \mathbf{G}_{cp,s} & \mathbf{G}_{cp,cp} \end{bmatrix} \mathbf{q}(t). \quad (7.24)$$

In the low mass case, the fluid coupling plays a mass-like role in its influence on the partition dynamics. However, the fluid also selectively couples neighbouring segments due to the Laplace operator underlying $-\mathbf{G}_{cp,cp}^{-1}$. Once again $\mathbf{G}_{cp,s}$ scales the input to each segment, but in this case $-\mathbf{G}_{cp,cp}^{-1}$ replaces \mathbf{M}^{-1} in \mathbf{B}_E . Assuming a rectangular fluid region once again, $(-\mathbf{G}_{cp,cp}^{-1}\mathbf{G}_{cp,s})_i \approx 0 \quad i \neq 1$, because the second derivative of $(\mathbf{G}_{cp,s})_i$ is zero away from the stapes boundary. Therefore, we have the required conditions for a travelling wave that starts strictly at the stapes and moves between neighbouring segments along the cochlear partition.

7.2.3.3 An Intermediate Value for $M(\xi_i)$

In this section, we have seen that the size of $M(\xi_i)$ determines whether the cochlear partition behaves like a set of isolated resonators, or if coupling effects are essential to its response. In practice, the cochlea cannot lie at either of the extremes that we consider, but rather at some intermediate value. Therefore, both terms in $[\mathbf{M} - \mathbf{G}_{cp,cp}]^{-1}$ will influence the partition behaviour. Figure 7.1 on the next page shows the effect of the fluid coupling for a range of intermediate values for $M(\xi_i)$. Section 7.3.2.3 shows that when $M(\xi_i)$ is determined by the volume and density of the cochlear partition, it lies close to the low mass end of the range.

7.2.4 Discussion

The mathematical manipulations presented in this section show clearly that otherwise isolated segments can be coupled through the cochlear fluid. This result is in agreement with previous analytical and numerical cochlear models, but the new contribution makes the action and nature of the coupling apparent. The dependence of the nature of the coupling on the segment mass shows that care must be taken when choosing the size of parameters, in particular the mass, to ensure key behaviour is correctly captured. Taken together with section 5.2, this is a very general result. \mathbf{M} in this work is equivalent to $\mathbf{C}_E\mathbf{B}_E$ in that section. The matrix $[\mathbf{I} - \mathbf{B}_E\mathbf{F}^{-1}\mathbf{C}_E]^{-1}$, and its equivalent $[\mathbf{I} + \mathbf{B}_E\mathbf{Q}[\mathbf{K} - \mathbf{R}\mathbf{C}_E\mathbf{B}_E\mathbf{Q}]^{-1}\mathbf{R}\mathbf{C}_E]$, determine the fluid coupling that the partition experiences. The second expression for the matrix allows us to be more specific. The identity matrix \mathbf{I} in that expression accounts for feedback through the segment model itself, and the term $\mathbf{Q}[\mathbf{K} - \mathbf{R}\mathbf{C}_E\mathbf{B}_E\mathbf{Q}]^{-1}\mathbf{R}$ accounts for feedback through the fluid. This justifies the name ‘‘fluid coupling matrix’’ that we




Figure 7.1: The effect of M on $[\mathbf{M} - \mathbf{G}_{cp,cp}]^{-1}$ (left) and $-\mathbf{G}_{cp,cp}^{-1} \mathbf{G}_{cp,s}$ (right). (Left) Values of M from left: 2.5×10^3 , 2.5 , 2.5×10^{-1} , 2.5×10^{-2} , 2.5×10^{-3} , $2.5 \times 10^{-4} \text{ g} \cdot \text{cm}^{-2}$. (Spatial arrangement for presentation only, $[\mathbf{M} - \mathbf{G}_{cp,cp}]^{-1}$ is insensitive to i away from boundaries.) The final starred trend is data from a 1D finite difference model for $M = 0$. This trend is normalised to 1.0, rather than by multiplying by M as in the case of the others. Note that as M is decreased, the coupling between segments increases. For low M , neighbouring segments are strongly coupled. (Right) Colours and M values match (left). Note that for small M , the range of segments that experience a significant stapes input is small. $M = 2.5 \times 10^{-3} \text{ g} \cdot \text{cm}^{-2}$ is a physiologically justifiable value for the mass of the partition (see section 7.3.2.3), and $M = 2.5 \times 10^3 \text{ g} \cdot \text{cm}^{-2}$ is sufficiently large to ensure that fluid coupling between the segments is small.

coined for it in section 5.3.1.3. The coupling through the fluid at any time step can be assessed by studying the linearisation of the radial segment models. (Terms that affect \mathbf{B}_E will influence the fluid coupling matrix.) Furthermore, it should extend naturally to two scalae cochlear models when a two scalae form of the monolithic system equation is used.

7.3 Steps to Build the Model

The mathematical analysis presented in section 7.2 provides a firm description of the coupling between radial segments. In terms of clarity, no numerical simulation can provide a better answer to the criticism outlined in section 7.1.3.4, so the mathematical analysis could easily “steal the show” in this chapter. However, we still need a simple example to demonstrate how the proposed framework may be used, and the numerical results that will be produced will confirm the correctness of our derivation, while offering another view of the result that is accessible to people who prefer a less mathematical approach.

In this section, we outline how the basic model required for this work is assembled using the framework described in chapter 4.

7.3.1 Select a Cochlear Partition Model Type

7.3.1.1 Model Inputs and Outputs

Our first task is to select a suitable cochlear partition dynamic equation. We only need a linear cochlear model in order to compare the model against [Bell \(2010\)](#), so equation (4.33) is suitable. The Q values that will be used in the cochlea are higher than the biological passive response, as they match the active cochlea. Therefore, the use of a linear model is a little contrived, but will be suitable for the purposes of this section.

We need to manipulate equation (4.33) to obtain a state space model of the form given in equations (4.18) and (4.19). To do this, we need to decide whether our fluid model should use pressure (4.14) or potential (4.16) as the solution variable. In both cases the fluid is treated as inviscid, irrotational, and incompressible. Therefore, the only energy storage element in the fluid is its mass. The kinetic energy per unit volume stored in the fluid is given by:

$$W = \frac{1}{2}\rho\mathbf{v}^2, \quad (7.25)$$

where W is the energy density at a particular point and time. The intensity \mathbf{I} (the vector flux of W) is defined by:

$$\mathbf{I} = p\mathbf{v}. \quad (7.26)$$

The two equations satisfy the requirements of conservation of energy, because $\partial W/\partial t = -\nabla \cdot \mathbf{I}$. (These relationships are specific to these fluid models. [Lighthill \(1978\)](#) section 1.3 discusses these relationships in the more general setting of linear acoustic waves.)

We choose the (2D) potential equation (4.16) as the fluid domain equation so that we can perform these energy calculations in the fluid without having to obtain the velocity field from a temporal integration. This also allows us to demonstrate the process of calculating the pressure difference across the partition as a data state. When expressed using the fluid potential, the equations become:

$$W = \frac{1}{2}\rho(\nabla\phi)^2, \quad (7.27)$$

$$\mathbf{I} = -\rho\left(\frac{\partial}{\partial t}\phi\right)\nabla\phi. \quad (7.28)$$

7.3.1.2 The Dynamic Model

Having selected the variables that the partition model should take as its input and provide as its output, we return to specifying the dynamic state space model. We recall equation (4.34), the equivalent

of (4.33) for potential flow, and convert the equation to its transfer function form:

$$\frac{v_{BM}(x, s)}{\Delta\phi_{BM}(x, s)} = \frac{\rho s^2}{M(x)s^2 + R(x)s + K(x)}, \quad (7.29)$$

where s is the Laplace transform variable. This is converted to a state space model. First, we convert the right-hand side to a proper fraction:

$$\frac{v_{BM}(x, s)}{\Delta\phi_{BM}(x, s)} = \frac{\rho}{M(x)} - \frac{\frac{\rho R(x)}{M(x)}s + \frac{\rho K(x)}{M(x)}}{M(x)s^2 + R(x)s + K(x)}. \quad (7.30)$$

Let $v_{BM}(x, s) = k_1\Delta\phi_{BM}(x, s) + k_2(x, s)$, where $k_1 = \rho/M(x)$ and k_2 are the single scala form of the constants required for the segment output linearisation discussed in section 6.2.2 and section 6.3.1. $k_2(x, s)/\Delta\phi_{BM}(x, s)$ is found from the strictly causal transfer function:

$$\frac{k_2(x, s)}{\Delta\phi_{BM}(x, s)} = -\frac{\frac{\rho R(x)}{M(x)}s + \frac{\rho K(x)}{M(x)}}{M(x)s^2 + R(x)s + K(x)} \times \frac{[w]}{[w]}, \quad (7.31)$$

where we introduce a signal $[w]/[w]$. This is one method of obtaining a state space model from a transfer function.

The denominator of the equation gives us our first state update equation:

$$\dot{X}_1 = s[sw] = -\frac{R(x)}{M(x)}[sw] - \frac{K(x)}{M(x)}[w] + \frac{1}{M(x)}\Delta\phi_{BM}(x, s), \quad (7.32)$$

where the $[\]$ indicate that the quantity is treated as a signal. The second state update equation is simply:

$$\dot{X}_2 = s[w] = [sw]. \quad (7.33)$$

These equations can be written in matrix notation as:

$$\frac{d}{dt} \begin{bmatrix} X_1 \\ X_2 \end{bmatrix} = \begin{bmatrix} -\frac{R(x)}{M(x)} & -\frac{K(x)}{M(x)} \\ 1 & 0 \end{bmatrix} \begin{bmatrix} X_1 \\ X_2 \end{bmatrix} + \begin{bmatrix} \frac{1}{M(x)} \\ 0 \end{bmatrix} \Delta\phi_{BM}(x, t). \quad (7.34)$$

This completes the state update equations required for the model.

$k_2(x, t)$ is found from the numerator of (7.31). The output equation is then:

$$k_1(x, t) = \frac{\rho}{M(x)}, \quad k_2(x, t) = \begin{bmatrix} -\frac{\rho R(x)}{M(x)} & -\frac{\rho K(x)}{M(x)} \end{bmatrix} \begin{bmatrix} X_1 \\ X_2 \end{bmatrix}. \quad (7.35)$$

This completes the dynamic state model.

7.3.1.3 The Data Model

We now consider what data outputs and data states that we might require in the model. In this case, these outputs will be used to monitor the performance of the system, rather than for a neural feedback loop. We already have the velocity of the partition, so we log that to file. The data output equation for the displacement of the partition, $d_{BM}(x, s)$, also depends on the dynamic states only:

$$d_{BM}(x, t) = \begin{bmatrix} \rho & 0 \end{bmatrix} \begin{bmatrix} X_1 \\ X_2 \end{bmatrix}. \quad (7.36)$$

(In this case, X_1 is the displacement scaled by the fluid density.) Using $v_{BM}(x, t)$ and $d_{BM}(x, t)$ we can compute the energy stored in the partition using:

$$E(x, t) = \frac{1}{2}M(x)v_{BM}^2(x, t) + \frac{1}{2}K(x)d_{BM}^2(x, t). \quad (7.37)$$

If both $v_{BM}(x, t)$ and $d_{BM}(x, t)$ are logged to file and $M(x)$ and $K(x)$ are known, we could also postprocess this value.

The acceleration of the partition is related to the velocity by $a_{BM}(x, s) = sv_{BM}(x, s)$. Comparing this to (7.29) shows that the transfer function $a_{BM}(x, s)/\Delta\phi_{BM}(x, t)$ is noncausal, and cannot be computed without introducing a nondominant pole. We elect not to compute $a_{BM}(x, s)$, because it is inconvenient on this occasion, and we will not require it.

The pressure difference across the cochlear partition is also a noncausal function:

$$p(x, s) = \rho s \Delta\phi_{BM}(x, s). \quad (7.38)$$

We convert this to a causal transfer function by inserting a nondominant pole at $\alpha s + 1$. Therefore, this data output requires one data state:

$$\frac{d}{dt}X_d = -\frac{1}{\alpha}X_d + \frac{1}{\alpha}\Delta\phi_{BM}(x, t), \quad (7.39)$$

and the output equation is:

$$p(x, t) = \frac{\rho}{\alpha}\Delta\phi_{BM}(x, t) - \frac{\rho}{\alpha}X_d. \quad (7.40)$$

Setting $\alpha = \Delta t/2$ produces an estimate that nominally has the same dynamics as the backward difference implementation of a derivative.

In this section, we have described the dynamic model and the data model required for the segment. The dynamic model takes $\Delta\phi_{BM}(x, t)$ as its input and provides $v_{BM}(x, t)$ as its output. This model requires two states. We will have four data outputs: $v_{BM}(x, t)$, $d_{BM}(x, t)$, $p(x, t)$, and $E(x, t)$. $v_{BM}(x, t)$ has already been computed, $d_{BM}(x, t)$ requires an extra data output equation only, $p(x, t)$ requires a

single state and data output equation, and $E(x, t)$ requires a nonlinear data output equation, based on the dynamic states.

7.3.1.4 Implementation

We implement this theory by extending the `OrganOfCortiComplexSegment` class to create a new class `PotentialSegment`. The `OrganOfCortiComplexSegment` class provides the majority of the functionality that we need, so all that is required is to setup the matrices for the dynamic and data state models. The class also allows meaningful names for the output variables to be provided that will be used when writing out data.

7.3.2 Select Tuning Parameters

The `PotentialSegment` class specification required three variables related to the partition: $M(x)$, $R(x)$, and $K(x)$. The fluid density, ρ , also appears as a parameter because it was introduced by the conversion from p to ϕ . The segment model is general because these have not been specified in that class. To actually make a segment, we need to provide this data, and an object inheriting from the `Parameters` class is used to do this.

7.3.2.1 Unit System

The most basic choice that must be made when specifying parameter values is what system of units to use. MKS (meter-kilogram-second) are the normal SI units. CGS (centimetre-gram-second) is also an SI unit system that is popular in cochlear modelling because the quantities relevant to cochlear models are well scaled in this system. The choice would often be swayed by the unit system adopted in the most relevant works to be compared against. In this case, the results presented in [Bell \(2010\)](#) do not make extensive use of one system or another, so we adopt the CGS system. In this system ρ is conveniently scaled to $1 \text{ g} \cdot \text{cm}^{-3}$.

The units of $M(x)$, $R(x)$, and $K(x)$ need to be considered. The general form of equation (4.33) has pressure $\Delta P_{BM}(x, t)$ with units $\text{dyn} \cdot \text{cm}^{-2}$ on the left-hand side⁸. Therefore, for consistency, the units of the mass, damping, and stiffness must be $\text{g} \cdot \text{cm}^{-2}$, $\text{dyn} \cdot \text{s} \cdot \text{cm}^{-3}$, and $\text{dyn} \cdot \text{s}^2 \cdot \text{cm}^{-3}$ respectively. The variable are expressed per unit area because the width and length of the basilar membrane have been cancelled to arrive at equation (4.33).

⁸ $1 \text{ dyn} \cdot \text{cm}^{-2} = 0.1 \text{ Pa}$, $1 \text{ dyn} = 1 \text{ g} \cdot \text{cm} \cdot \text{s}^{-2}$.

7.3.2.2 Cochlear Map and Quality Factor, Q

Bell (2010) adopts the cochlear mapping of Greenwood (1961) (see figure 3.11):

$$f(\tilde{x}) = 165.4 \left(10^{0.6\tilde{x}} - 1 \right), \quad (7.41)$$

where \tilde{x} [cm] is the distance along the cochlear partition from the apex of the cochlea. In deriving this mapping, Greenwood assumed a cochlear length of 3.5 cm for the human cochlea partition. We will also use this value when the need arises in section 7.3.3.

Bell draws Q values from Shera *et al.* (2002) that presented Q_{ERB} values (Q values based on an equivalent rectangular bandwidth). The equation for Q_{ERB} presented in Shera *et al.* (2002) and used in Bell (2010) is:

$$Q_{ERB}(f) = 12.7 (f/1000)^{0.3}. \quad (7.42)$$

The units of f in this function are kilohertz, and it is applicable at frequencies of 1 kHz and above according to Shera *et al.* (2002). Equations (7.41) and (7.42) therefore provide the frequency map and quality factor that we will try to replicate in our parameter selection.

7.3.2.3 A Physiological Value for $M(x)$

The mass of the partition varies less than its stiffness, and it is often treated as constant along the length of the partition. We assume that $M(x) = M_1$ to first approximation in this investigation.

Neely (1981a) table I compares the mass values that had been used in five papers published before 1980. Four of the five lie in the range $0.1 - 0.2 \text{ g} \cdot \text{cm}^{-2}$, the remaining value is much lower, at $0.01 \text{ g} \cdot \text{cm}^{-2}$. In his PhD thesis (Neely 1981b), Neely used anatomical data for the width of the basilar membrane and the fluid density ρ to estimate the mass of the basilar membrane. The resulting values lay in the range 4.5×10^{-3} to $2.7 \times 10^{-2} \text{ g} \cdot \text{cm}^{-2}$, lower overall than for previous models. Mammano and Nobili (1993) also computed the mass of the partition by assuming it to have a density similar to the cochlear fluid.

In this work, we assume $M_1 = 2.5 \times 10^{-3} \text{ g} \cdot \text{cm}^{-2}$, based on the fluid density and an effective height h for the cochlear partition. This value is compatible with the height of the cochlear partition (estimated at $5 \times 10^{-3} \text{ cm}$ from data in Wever 1949). de La Rochefoucauld and Olson (2007) note that h may be less than the actual height.

When comparing mass values in the literature, the following points are important to note: the order of magnitude of $M(x)$ has decreased significantly between early models and more recent models; there are significant differences between species; and modern models include detailed micromechanical models, so none of the parameters in these models may be directly comparable to $M(x)$.

The results for $M_1 = 2.5 \times 10^{-3} \text{ g} \cdot \text{cm}^{-2}$ will be compared to the response of a system with significantly higher mass $M_2 = 2.5 \times 10^{+3} \text{ g} \cdot \text{cm}^{-2}$. Section 7.2 shows that the partition responds like a set of isolated resonators for this large mass value, allowing us to easily compare and contrast the results for M_1 to the descriptions of the “apparent” travelling wave.

7.3.2.4 Computing $K(x)$ and $R(x)$

For an isolated second order mass-spring-damper system, $K(x)$ can be computed directly from equation (7.1), together with the cochlear frequency map (7.41) and the selected value of $M(x)$. (Strictly speaking this neglects the effect that damping has on the resonance frequency, however $Q > 12.7$ over the frequency range of interest, making the correction negligible.)

With the values of $Q(x)$, $M(x)$, and $K(x)$ specified, we can compute $R(x)$ using the formula:

$$Q(x) = \frac{\sqrt{M(x)K(x)}}{R(x)}. \quad (7.43)$$

This formula technically uses Q_{3dB} rather than Q_{ERB} , but Bell (2010) and Shera *et al.* (2010) treat them as equivalent. (This is fair because Q_{ERB} aims to approximate Q_{3dB} and the deviation is slight for most systems.)

These expressions can be used for the M_2 case, but section 7.2 shows that they are not justified for the value of M_1 selected above, because the segments do not act in isolation, but rather influence each other strongly. For simplicity, we replace $M(x)$ in (7.1) by $M_{eff} = 18 \times 10^{-3} \text{ g} \cdot \text{cm}^{-2}$, a value selected by hand to provide reasonable agreement between the cochlear map and velocity waveform at 2844 Hz (with a maximal response at 1.4 cm)⁹. Figure 7.2 shows the match attained.

7.3.2.5 Implementation

We have now specified spatial functions for $M(x)$, $R(x)$, and $K(x)$ that are based on the same physiological data as Bell (2010), apart from ρ and M that we supply from the literature. The `PotentialSegment` requires this data to be provided by an object inherited from the `Parameters` class. The `Parameters` class imposes minimal restrictions on the interface of our new class, which we will name `Bell2010Parameters` as a mnemonic. All it requires is that our class provide a method with the declaration:

```
virtual void calculate_values (double x_location);
```

⁹The effect of this change is to raise the resonance frequency of the segment significantly above the frequency required for the cochlear map. In this case, the local resonance does not determine the resulting tuning. de La Rochefoucauld and Olson (2007) found that local resonance is important in determining the cochlear map at the base of the gerbil cochlea, but has diminished importance at more apical locations.




Figure 7.2: Velocity response of the cochlear partition to the three input frequencies indicated. (Left) The response for a high mass scenario $M(x) = 2.5 \times 10^3 \text{ g}\cdot\text{cm}^{-2}$. Note the narrow, symmetrical response. (Right) The response for a low mass scenario $M(x) = 2.5 \times 10^{-3} \text{ g}\cdot\text{cm}^{-2}$. The response is asymmetrical and broader, despite a comparable Q for individual segments. Active mechanisms would sharpen the response. This figure presents frequency domain results, computed using the model’s state space matrices and equation (45) in Elliott *et al.* (2007).

This method is called by the `OrganOfCortiComplexCompilation` to initialise the `Parameters` object before it is passed to the segment. We will use this method to update the values of four variables, named `mass`, `damping`, `stiffness`, and `density` for clarity, using the relationships already discussed.

We have control over the way that `PotentialSegment` calls for data from `Bell2010Parameters`, so we can define the interface between these two classes. It is sensible to use a function like:

```
std::vector <double> get_parameters ();
```

because we want to return four double precision values. We order these values in the `std::vector` using the same sequence as above.

The constructor for the class can take a list of arguments. We can use this to allow M to vary, either by providing a boolean input selecting M_1 or M_2 , or by allowing values for M and M_{eff} to be supplied by the user.

7.3.3 Specify a Fluid Region

7.3.3.1 Dimensions

In order to specify the fluid model, we require the set of partial differential equations to be solved on the domain, the geometry of the region, and boundary conditions to be imposed on the model. The differential equation has already been selected in section (7.3.1.1) based on a consideration of the data we want to view for the fluid. The boundary conditions will be discussed in the next section, which brings us to the geometry.

We will use a 2D fluid region. Equation (7.41) assumed a cochlear partition length of 3.5 cm. The frequency map in (7.41) has the special property that it minimises apical reflections (Puria and Allen 1991). Apical reflections were discussed in sections 4.2.5 and 4.6.2. This property would be lost if we simulate a truncated length of the cochlea. Bell (2010) focused on a frequency range of 1 to 10 kHz because equation (7.42) from Shera *et al.* (2002) is only valid for frequencies above 1 kHz. According to equation (7.42), as $f \rightarrow 0$, $Q \rightarrow 0$ as well. This might create problems in equation (7.43) because $R \propto 1/Q$, but the overall dependencies of $Q(x)$ and $K(x)$ on frequency results in $R(x) \propto f^{0.7}(x)$. Therefore, provided we take care not to introduce numerical errors in the implementation, we can use the method outlined in section 7.3.2 below 1 kHz, but we must not rely on low frequency inputs or the response at the apex of the cochlea for our conclusions.

We need to supply a value for the height of the cochlea. We select the value 0.1 cm, which was discussed in section 4.2.5. For the majority of the results we will use a single scala cochlear model. In the single scalae model, the cochlear fluid was divided into 750 cells along the cochlear partition (for 1500 segments) and 30 cells in the height direction. Some results from an equivalent two scalae model will be used to show that our implementation of the mechanisms for two scalae cochleae are correct, and that our single scala results extend to two scalae models. In this model, a mesh with 1120×32 cells per scalae (2240 segments) was used.

7.3.3.2 Postprocessing the Fluid Solution

Section 7.3.1.1 discussed the equations required to compute the energy density and intensity in the cochlea fluid. They are functions of velocity and pressure. Discussing how these equations may be computed provides us an opportunity to consider the pros and cons of postprocessing data using the deal.II library versus a subsequent tool.

The velocity field at a given time can be computed from the potential field at that time, as $\mathbf{v} = -\nabla\phi$. This can be done either in the deal.II simulation or as a postprocessing step. Computing the gradient in deal.II is more accurate, because the process of outputting data to another program can introduce errors. On the other hand, leaving the postprocessing until the data analysis stage avoids computing and storing values for the velocity that may never be used.

To compute the pressure, we can use the backward difference formula:

$$p_n(x, y) = \frac{\rho}{\Delta t} (\phi_n(x, y) - \phi_{n-1}(x, y)), \quad (7.44)$$

where n refers to the solution just computed and $n - 1$ the previous solution, two sequential time steps. We generally don't need to see the result at each simulation time step, so we are unlikely to be writing sequential time steps to file otherwise. Therefore, whether we do this computation in deal.II or in a postprocessing program we need to add another spatial vector to the output file (either

p_n or ϕ_{n-1}). A set of state space models, like the one discussed in section 7.3.1.3, could be used to compute the pressure, but a backward difference gives competitive performance and is simple. If we did use state space models, we would need to do the computation in deal.II, because the amount of data that would need to be written out to do this as a postprocessing stage is excessive.

There are times when a computation must be, or is best, done using the deal.II library and others when a postprocessing program can do the same job. Sometimes performing the postprocessing up front reduces the data written out, but often it means that valuable computation time is spent on results that might never be used. For these investigations, both \mathbf{v} and p were computed in deal.II. This is because we also log the time data for a small set of points. W and \mathbf{I} were computed in the postprocessing and visualisation package Paraview (Squillacote 2007).

7.3.4 Choose Boundary Conditions

7.3.4.1 Boundary Indicators

We now consider the boundary conditions that must be applied to complete the description of the fluid domain. The cochlear partition boundary condition will be provided by the `PotentialSegment` class, and we simply need to indicate where this boundary is located. To do this, we mark the finite element cells above the basilar membrane interface with an integer “material index”. We indicate each of the remaining boundaries using a unique “boundary indicator”. These are variables provided by deal.II for these types of situations.

7.3.4.2 The Cochlear Wall

The boney wall of the cochlea is the simplest boundary to deal with. On this boundary, we want to enforce zero normal flow. This is the “default” condition obtained if no other boundary condition is enforced, so we simply ignore faces marked as bone.

7.3.4.3 The Helicotrema and Round Window

Next we consider the round window and helicotrema boundary conditions. Section 7.3.3 discussed the fact the cochlear map already minimises apical reflections, so applying a zero Dirichlet condition on the pressure, and hence ϕ , at the helicotrema will not cause problems. For simplicity we adopt the boundary condition $\phi = 0$ on Γ_h .

7.3.4.4 The Stapes

The stapes is the final boundary condition to consider. Section 4.5.1 discusses the importance of the stapes model, especially for nonlinear cochlear models. A simple stapes, or rather middle ear, model is:

$$v_S(s) = \frac{s}{s^2 M_m + s R_m + K_m} \left(\frac{A_m}{G_m A_S} P_e(s) - \bar{P}_S(s) \right), \quad (7.45)$$

where M_m , R_m , and K_m are lumped mass, damping, and stiffness terms for the middle ear, A_m and A_S are the effective areas of the eardrum and stapes respectively, and G_m is the middle ear gain. P_e is the pressure in the ear canal, and \bar{P}_S the average fluid pressure acting on the stapes. (This is essentially the middle ear model used in [Neely and Kim \(1986\)](#), converted to the time domain.)

Equation (7.45) shows that the back pressure of the fluid acting on the stapes \bar{P}_S affects the velocity of the stapes. This is to be expected, but \bar{P}_S is dependent on the response of all the radial segments, and also has an effect on their responses. In section 7.2, we removed the coupling via the stapes model to clarify the segment to segment coupling. Equation (7.45) shows that it is necessary to do this again. Therefore, we will enforce a velocity at the stapes as the input to the model. This is acceptable because we are studying linear systems.

7.3.5 Assemble the Pieces

This framework follows a modular approach. The `PotentialSegment` and `Bell2010Parameters` classes can be tested in isolation, and together, before being used in a cochlear model. Similarly, the fluid model can be tested for a variety of input conditions without using any radial segments. Simple tests on the classes can quickly determine any obvious mistakes in their implementation. Once they are behaving as expected in isolation (and this expected behaviour is well defined), it is straightforward to combine them.

Assume that we have a working fluid model (say from a previous experiment). This fluid model uses an `OrganOfCortiComplexCompilation` that in turn requires `OrganOfCortiComplexSegment` and `Parameters` objects. To change the organ of Corti boundary condition, we simply supply the new class names to the `OrganOfCortiComplexCompilation` and check some other selections such as the material index, whether to use two scalae or just one, and the optional parameters list for the new class derived from `Parameters`. If there are any `OtherBoundary` class objects, then similar changes are required to include (or remove) these. Finally, we need to specify how to obtain the fluid geometry (supplying either a file or the commands to create it in the simulation). These few changes convert an existing cochlear model into a new model, and we do not expect to do much debugging of the new model, because we have reused tried and tested code.




Figure 7.3: Steady state (settled) response to a sine wave. Stapes input velocity: 1.0 cms^{-1} at 2844 Hz, with a 4 ms raised cosine onset. (Left) High mass, M_2 . (Right) Low mass, M_1 . Four time instants were selected to make the envelope of the response clear. For comparison, see figure 7.2.

7.4 Results

In this section, results from the model discussed above are presented. We first look at the steady state behaviour of the system to show that the time domain results agree with the frequency domain results already presented. The transient response of the system is then studied to highlight the qualitative differences between the behaviour for the low mass M_1 and high mass M_2 cases. The transient response also shows the qualitative *similarity* between the M_2 case and the response of isolated segments. A major distinction between an apparent travelling wave and a true travelling wave is that the apparent travelling wave does not carry any energy. Section 7.4.3 presents postprocessed data concerning the energy flow in the cochlea, allowing us to discuss this distinction. Finally, we present some results from a two scalae model with (fairly contrived) boundary conditions that should result in identical results to the M_1 case. By using boundary conditions that should produce a known result, we verify that the mechanisms to incorporate two scalae in the model are correctly implemented. The application of the results to the question outlined will be discussed in the next section.

7.4.1 Steady State Behaviour

Figure 7.3 shows the response of the model to a single input frequency for two different mass conditions. The *only* difference between the responses is the change in parameters of the model. In particular, the fluid model is identical in both cases. The results show good agreement with figure 7.2 on page 159, confirming that both conditions are correctly implemented. See the figure caption for further details.




Figure 7.4: The transient response of isolated segments to a 2844 Hz sine wave. Parameters for the isolated segments are as for the (a) 0.7 cm, (b) 1.4 cm, and (c) 2.1 cm locations in the cochlear map.

7.4.2 Transient Behaviour

In this section we present five figures: three with time on the abscissa; and two with distance along the partition on the abscissa. Figure 7.4 on the current page shows the response of three isolated segments to a 2844 Hz input sine wave that represents a pressure difference across the segment. The segments are tuned using the cochlear frequency map and Q for the 0.7 cm, 1.4 cm, and 2.1 cm locations. These locations are the best place for 7751 Hz, 2844 Hz, and 979 Hz respectively, and will also be studied in the cochlear model.

Figures 7.5 and 7.6 show the response for the three locations listed above for a 2844 Hz sine wave velocity at the stapes. The results correspond to the M_2 and M_1 mass conditions respectively. For the high mass case M_2 , the response of each segment has the same shape as the corresponding uncoupled segment, with only the amplitude of the response changed. This is consistent with our expectations based on the mathematical results presented in section 7.2. The result confirms that when M is large the cochlear model behaves like a set of independent resonators.

On the other hand, the response for the M_2 condition, shown in figure 7.6, shows significant qualitative differences from the isolated segments. Firstly, all segments show a delay “ringing up” whereas only the 1.4 cm location (that corresponds to the input frequency, 2844 Hz) showed behaviour like this in the other results. The delay is not consistent with Q cycles of the 2844 Hz input – the M_1 response completes fewer sign changes in the ring up phase, despite the individual segments having




Figure 7.5: The transient response of three segments in the coupled cochlear model for the high mass condition, M_2 . The segments selected are located at (a) 0.7 cm, (b) 1.4 cm, and (c) 2.1 cm. Notice the close correspondence to figure 7.4. The two figures only differ by scaling factors.




Figure 7.6: The transient response of three segments in the coupled cochlear model for the low mass condition, M_1 . The segments selected are located at (a) 0.7 cm, (b) 1.4 cm, and (c) 2.1 cm. The response of all of the segments differs qualitatively from that seen in figures 7.4 and 7.5.

the same Q . Perhaps most significantly, the 2.1 cm location (979 Hz best frequency) shows an initial response that dies away completely after approximately 30 ms. The onset of the sinewave evokes a response at all locations on the basilar membrane, leading to the initial response. The steady state response is consistent with a segment apical of the best place for the input frequency, which does not receive significant input from a true travelling wave.

The next two figures present the velocity response of the cochlear partition to an impulse at the stapes for five time instants. Figures 7.7 and 7.8 below show results for the high mass case M_2 and the low mass case M_1 respectively. Both results show what appears to be a travelling wave, and the speed of the “wave” is similar. In the high mass case, a large proportion of the partition is involved in the response (the result at 2.5 ms is a particularly good example). Over the 20 ms presented, the peak value of the response decayed by three orders of magnitude, from $\approx 3 \times 10^{-4} \text{ cm} \cdot \text{s}$ to $\approx 3 \times 10^{-7} \text{ cm} \cdot \text{s}^{-1}$. Therefore, the behaviour can be interpreted as the result of the impulse exciting the entire partition, with the oscillations dying away more rapidly in the basal region than the apical region. This interpretation is consistent with the segments behaving as if they are isolated, resulting in an apparent travelling wave.

On the other hand, in the low mass case, the response is always restricted to a narrow band on the cochlear partition. The amplitude of the velocity remains fairly constant, decreasing from $\approx 8 \times 10^{-2} \text{ cm} \cdot \text{s}^{-1}$ to $\approx 1.5 \times 10^{-2} \text{ cm} \cdot \text{s}^{-1}$. The fact that the response at any instant is restricted to a narrow band is consistent with a true travelling wave, where the signal progresses apically in the cochlea. In the high mass case, with the large variation in the amplitude of the velocity, an early response in the apex of the cochlea might be masked by the larger response at the base (figures 7.7 and 7.5 suggest that this is occurring). However, in the low mass case the variation in the velocity amplitude is too small to mask an apical response in this manner.

7.4.3 Postprocessing and Visualisation

Section 7.3.1.1 discusses two measures related to the energy in the cochlear fluid that may be post-processed from the fluid domain solution: W , defined in equation (7.27); and \mathbf{I} , defined in equation (7.28). The energy stored by the cochlear partition E is defined in equation (7.37). E is an energy density measure with units $\text{g} \cdot \text{s}^{-2}$ because $M(x)$ and $K(x)$ are per unit area values¹⁰. Once again, the stapes input velocity was $1.0 \text{ cm} \cdot \text{s}^{-1}$ at 2844 Hz, with a 4 ms raised cosine onset.

Figure 7.9 on page 169 shows the fields W and \mathbf{I} in the basal half of the cochlea (0–1.75 cm from the stapes) for the high mass case M_2 . The results in all the figures are from the 22.5 ms time instant. This time instant was selected because the stapes input velocity is small, allowing the effect of the partition on the fluid to be seen. In the simulation with M_2 , the fluid motion due to the deformation of the partition is much smaller than the fluid flow directly from the stapes to the helicotrema and

¹⁰The usual unit for energy in the CGS system is the erg. $1 \text{ erg} = 1 \text{ g} \cdot \text{cm}^2 \cdot \text{s}^{-2} = 1 \times 10^{-7} \text{ J}$.




Figure 7.7: The velocity of the cochlear partition at five time instants for the high mass parameters, M_2 , in response to an impulse at the stapes. Based on the mathematical arguments and numerical results presented, this is the “apparent” travelling wave. The amplitude of the wave varies over three orders of magnitude as time progresses (see exponents on the y-axis).

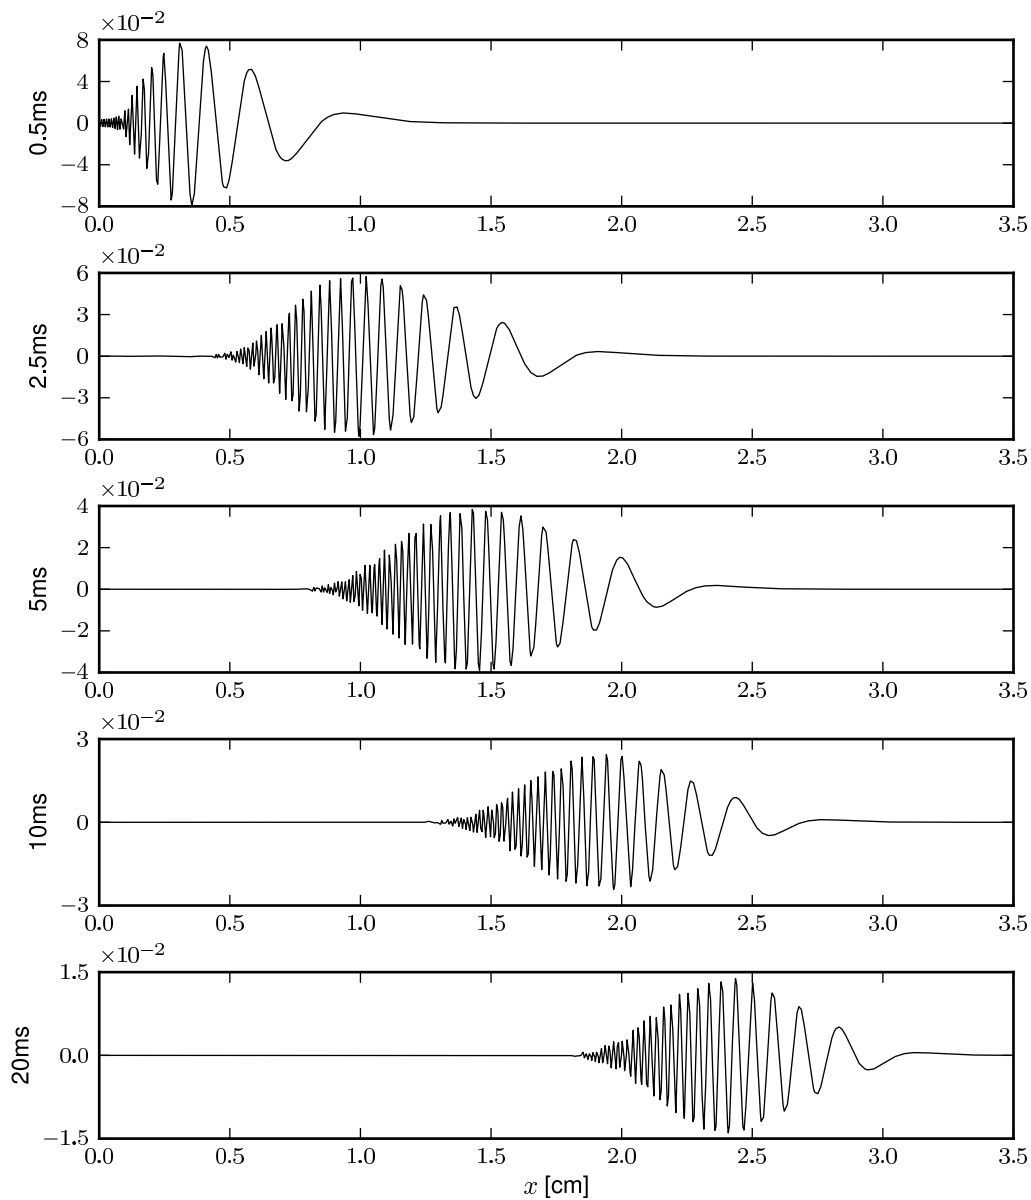


Figure 7.8: The velocity of the cochlear partition at five time instants for the low mass parameters, M_1 , in response to an impulse at the stapes. Based on the mathematical arguments and numerical results presented, this is the travelling wave usually attributed to the cochlea.




Figure 7.9: Postprocessed fluid fields and cochlear partition data for the high mass case, M_2 . (a) Energy density W (see equation 7.27). (b) Intensity I (see equation 7.28). Colour scale relative to magnitude in each panel. Deep blue, minimum value; red, maximum value. Arrows in (b) indicate direction of I . (c) Plot of partition energy density E along the cochlea (see equation 7.37).

back. (This is consistent with the physical interpretation that the partition is modelled as heavy and stiff, so it deforms very little.) The time instant was also selected to highlight an occasion when the energy flow (intensity) is back towards the stapes and *out* of the the cochlea. The response shows that just a narrow band of the cochlear partition, centred on 1.4 cm, stores significant energy in this case. Similarly, the fluid and the structure only interact strongly in this region.

Figure 7.10 on the next page shows the low mass case M_1 , once again for the 22.5 ms time instant. In this case, an extended length of the fluid-partition interface is involved in the response. Comparing panel (b) and (c) in particular shows the close interaction between the fluid and structure, with the intensity field indicating that the major energy flow is from one point on the cochlear partition, through the fluid, and to another slightly apical point on the partition. Through this process, the position of peaks and troughs in panel (c) shifts apically with time, although the envelope of the response does not change. In this case, the direction of the intensity is *always* from stapes to apex.

The intensity in the vicinity of the partition contains a component normal to the partition that describes energy exchange between the fluid and the structure, and a component tangential to the partition that represents energy transfer towards the apex of the cochlea. Figure 7.11 on the following page shows just the normal (\hat{j}) component of I in panel (a), allowing us to compare the energy exchange profile to the energy stored by the partition. From the figure, it is clear that the partition transfers energy to the fluid on the positive slope of the $E(x)$ graph, and takes up energy from the




Figure 7.10: Postprocessed fluid fields and cochlear partition data for the low mass case, M_1 . (a) Energy density W (see equation 7.27). (b) Intensity \mathbf{I} (see equation 7.28). Colour scale relative to magnitude in each panel. Deep blue, minimum value; red, maximum value. Arrows in (b) indicate direction of \mathbf{I} . (c) Plot of partition energy density E along the cochlea (see equation 7.37).




Figure 7.11: Postprocessed fluid fields and cochlear partition data for the low mass case, M_1 . (a) Component of the intensity \mathbf{I} in the height direction \hat{j} . Colour scale relative to magnitude in each panel. Deep blue, minimum value; red, maximum value. Arrows indicate direction of \mathbf{I} . (b) Plot of partition energy density E along the cochlea (see equation 7.37).




Figure 7.12: Steady state (settled) response to a sine wave in a two scalae cochlear model. Stapes input velocity: 1.0 cm s^{-1} at 2844 Hz, with a 4 ms raised cosine onset. Four time instants were selected to make the envelope of the response clear. For comparison see figures 7.2 and 7.3.

fluid on the negative slope of that graph. (This can be seen most clearly in the region around the 1.0 cm point, before the wave length of the response becomes too short.) The action of this energy transfer would be to shift the peak of the energy distribution slightly apically in the next time instant. Interestingly, maximums and minimums of $E(x)$ correspond to zero crossings in $\mathbf{I} \cdot \hat{\mathbf{j}}$, showing that all energy transfer occurs at points where $E(x)$ is varying.

7.4.4 Results from a Two Scalae Model

In this section, we present results to confirm that the mechanisms which facilitate two scalae models are correctly implemented. The fluid domain for the two scalae model consisted of two 3.5 by 0.1 cm channels, separated by a distance of $10 \mu\text{m}$ to allow space for the helper cells in the cochlear partition (see section 6.2.3.1). $\phi = 0$ was enforced on the apical boundary of both scale, which is the equivalent of the Γ_h used previously. The basal boundary in the scala tympani was driven in antiphase to the scala vestibuli stapes boundary. Once again, the stapes input velocity was a 1.0 cm s^{-1} , 2844 Hz sine wave, with a 4 ms raised cosine onset. Only results from the M_1 case are presented.

Figure 7.12 shows the settled response of the model for four time instants. As expected, the envelope of the response shows good agreement with figure 7.3. Slightly different time instants were selected for the figure. This was necessary because the results shown are from a slightly different point on the cochlear partition – a result of the different mesh refinement between the two scalae and single scala models. The phase of the response is particularly sensitive to this change, and new time points were selected from the (fairly sparse) set of spatial time steps written to file in order to capture the

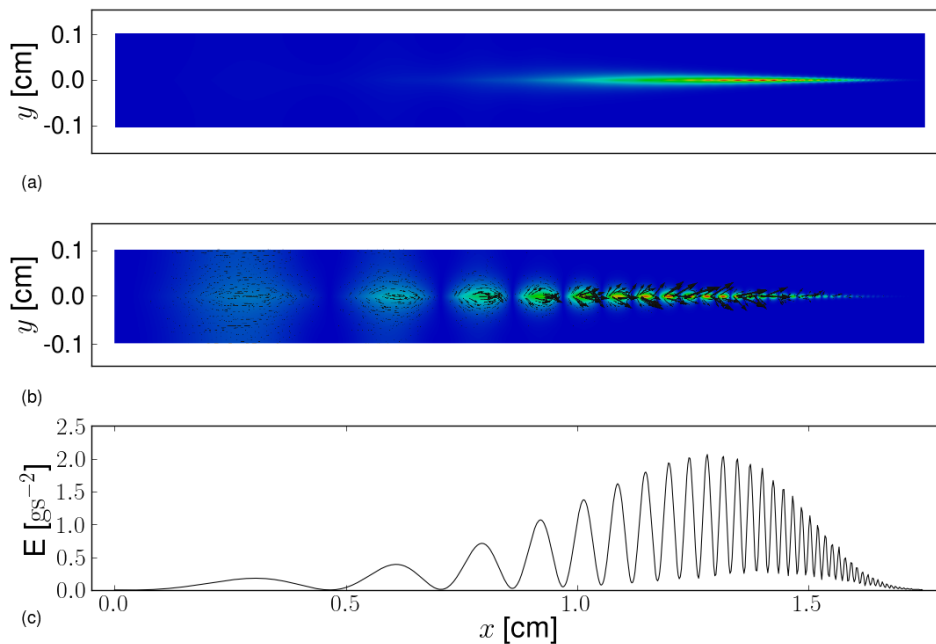


Figure 7.13: Postprocessed fluid fields and cochlear partition data for the two scalae low mass case, M_1 . (a) Energy density W (see equation 7.27). (b) Intensity \mathbf{I} (see equation 7.28). Colour scale relative to magnitude in each panel. Deep blue, minimum value; red, maximum value. Arrows in (b) indicate direction of \mathbf{I} . (c) Plot of partition energy density E along the cochlea (see equation 7.37). Note the close similarity to figure 7.10.

envelope again.

Figure 7.13 on the current page is equivalent to figure 7.10, and comparison of the plots of E in particular show the good agreement attained between the two scalae and the single scalae models. (Once again, the slight differences in the detail can be attributed to the slight difference in the mesh refinement.)

7.5 Discussion

Both the mathematical analysis and the numerical results show that two distinct types of behaviour can be obtained from the same mathematical model of the cochlea, depending on the size of the partition mass. If the mass is based on the dimensions of the partition and a density similar to the fluid density, then the resulting response shows a true travelling wave. We claim that this is a true travelling wave for two reasons: it carries energy; and the fluid supplies coupling between neighbouring segments. $M(x) = M_1 = 2.5 \times 10^{-3} \text{ g}\cdot\text{cm}^{-2}$ was used to study this condition. If the mass is significantly higher, the fluid coupling between segments has a negligible effect on their response, and they behave like isolated segments. The numerical results confirmed that when

$M(x) = M_2 = 2.5 \times 10^3 \text{ g}\cdot\text{cm}^{-2}$ the qualitative behaviour of the segments is the same as isolated segments with the same tuning.

Charles Babbs recently analysed a Helmholtz-Guyton cochlear model and observed an apparent travelling wave (Babbs 2011). Babbs cites work of Helmholtz that suggested the basilar membrane provided stiffness, while the fluid above and below the partition contributed mass. To the extent that this work shows the fluid provides significant loading of the partition by the fluid, this work agrees with Helmholtz' analysis, but Helmholtz did not anticipate the strong coupling between the segments that the fluid provides. Babbs continues by explaining that Guyton's contribution to the model was to specify that the "U" shaped fluid column between the stapes, a point on the partition, and the round window served as the fluid mass loading the partition – implying that there is no fluid coupling between the segments. With only the fluid contributing the mass, $M(x) \approx 0$. We consider this extreme case in figure 7.1, and the M_1 responses should also be representative. Therefore, our results show that this assumption is incorrect. In the low mass case, coupling between segments *cannot* be neglected. The computations that Babbs uses to show an apparent travelling wave on the cochlear partition miss this important detail. However, as we noted in section 7.1.3.2, Babbs also provided an interesting analysis of the stiffness and damping present in the basilar membrane. His analysis on this subject seems to be in agreement with other works on the topic.

Andrew Bell also recently published a paper discussing an apparent travelling wave (Bell 2010). He estimates the "delay" and "wave speed" of an apparent travelling wave that occurs in a graded set of independently excited resonators, using generic relationships between Q and the number of cycles an oscillator takes to "ring up". He finds that the predicted wave speed is comparable to physiological measurements, and suggests that this is evidence that an apparent travelling wave might be a consistent explanation of the available results. The values Bell adopted for Q were computed based on the delay of otoacoustic emissions in humans and other species (Shera *et al.* 2002). When Bell uses Q and frequency to compute travelling wave delay and ultimately speed, he is essentially reversing the computation done by Shera *et al.* Therefore, the similarity between the wave speed Bell calculated and physiological results seems to be inevitable from the method used to process the data.

The M_2 condition behaves like a set of isolated segments, so its behaviour can be used to study the apparent travelling wave in more detail. To this end, we adopted the cochlear map and Q values used by Bell. The results confirmed that the wave speed of the apparent wave is similar to the travelling wave observed for mass M_1 . A narrow region of the partition responds to a sine wave input in the M_2 case. This is seen as an advantage by Bell (personal communication), in particular because it provides an obvious link between place and frequency. We note that having a wide region of the partition involved in the response can be an advantage, if each hair cell contributes energy to amplify the input signal: the work is then shared between more hair cells, which reduces the amount of energy that each hair cell must supply. We found that the energy flow in the cochlear fluid (the

intensity) can sometimes be directed out of the cochlea when the segments act in isolation – this never occurs when a true travelling wave is present on the cochlear partition. Lynch III *et al.* (1982) found that the cochlear partition acts like a resistive load in the cochlear input impedance, agreeing with the true travelling wave scenario, but this does not rule out a small outward component for very small input velocities.

Bell's paper aimed to provide support for his surface acoustic wave cochlear model. The input to Bell's model is the absolute pressure in the cochlear fluid, detected by outer hair cells. (Bell assumes that outer hair cells are able to do this.) He assumes that the absolute pressure does not do any significant work on the cochlear partition, it only provides a signal. The output is a variation in the length of the outer hair cells. Equation (4.14) shows that the *acceleration* of the fluid-hair cell boundary must serve as the feedback to the fluid model. Depending on details of the relationship between pressure and acceleration, which is governed by the outer hair cells at low input levels, the feedback could show a high effective mass. As such, the high mass required for the partition to behave like a set of isolated segments does not inherently disprove Bell's model. However, if the density of the partition is *not* similar to water then there will be reflections at the fluid-partition boundary, which Bell does not consider. On the other hand, if the density *is* similar to water, then the conditions for a travelling wave on the cochlear partition (low mass and a stiffness element) are met, and the coupling between any apparent travelling wave in the vicinity of the hair cells and the true travelling wave due to the basilar membrane stiffness must be considered. Both waves rely on pressure as the input and acceleration as the output, so they will be acoustically coupled even if the structures between the hair cells and the basilar membrane do not provide physical coupling.

This work has provided some useful insights into the behaviour of a graded set of oscillators coupled by a fluid, but a number of limitations on the results have been noticed:

- Although high Q values have been used, which would imply an active cochlea, no sharpening mechanisms have been included in the M_1 case so the response remains broad;
- Bell's model takes the absolute pressure as an input, and should be studied with a two scalae model. Therefore, the results from the M_2 case can only provide information about sets of generic oscillators, rather than the specific interactions that might be present in Bell's proposed model.

Further work using this framework could remedy both of these shortcomings. A two scalae model should be used, allowing absolute pressures to be used as inputs when required. A two degree of freedom model, such as that proposed in Neely and Kim (1986) (and studied using this framework in Rapson *et al.* (2012) and Rapson and Tapson (2011b), see chapter 5), would provide scope for both the sharpening mechanisms required in the travelling wave theory, and separate degrees of freedom to represent the overall cochlear partition dynamics and the local hair cell dynamics. The tests on the two scalae model presented in this chapter are a first step towards this more detailed model.

7.6 Summary

This chapter has addressed two separate aims. Firstly, it has provided a detailed example of how the proposed framework can be used to build a cochlear model. The process was described in section 7.3, and examples of the type of results that can be obtained are provided in section 7.4.

The second aim of the chapter has been to acknowledge the presence of models that do not rely on a travelling wave in the cochlea, and to consider some of their claims. We set the scene by considering the modes in which technology projects and scientific research can fail. Simply, the diversity of models that have been proposed will ensure that some models will be wrong in part or totality. Given that the research has not converged on a single model, it is likely that even the best models available are only correct in part. Section 7.1.3 aims to provide a balanced overview of the more common criticisms of the travelling wave hypothesis. The analysis of one criticism was developed in more detail. The mathematical analysis in section 7.2 provides clear proof that the cochlear fluid can couple neighbouring segments, supporting a true travelling wave. This result was borne out by the numerical analysis that followed. The combined weight of evidence disproves key assumptions in the Helmholtz-Guyton model discussed by Babbs, because the model used is directly comparable to that model. The results can only contribute to the discussion on Bell's surface acoustic wave model because, although the models have important similarities, there are also differences to consider.

The analysis presented in section 7.2 also completes our study of the fluid-structure interaction in the cochlea. It shows that whether the ratio of segment mass to fluid density is high or low, the fluid-structure coupling in the cochlea is dominated by feedback from the neighbourhood of each segment. The fluid coupling matrix only depends on the geometry of the cochlea, the fluid density, and the segment mass. Therefore, the study considered all the pertinent variables. Section 5.2 showed that the fluid coupling matrix in nonlinear partition models have the same structure as in linear models. One specific term (\mathbf{B}_E) in the linearisation has an influence on the fluid coupling matrix. As the Laplace equations are a particularly accurate representation of the cochlear fluid over short time scales and small distances, the same fluid coupling matrix will characterise the behaviour of cochlear models in general. This result means that the mechanisms we have used to ensure that the fluid-structure interaction is captured correctly in the framework will extend to a wide class of cochlear models.

Chapter 8

Hopf, Van der Pol, and Silicon Cochleae

The original goal of this thesis was to develop a numerical cochlea model as a parallel track to work by our collaborators (André van Schaik and Tara Hamilton) on VLSI (Very Large Scale Integration) analogue silicon cochlea models, to allow comparison between the modelling approaches. My work with Dr Hamilton in 2010 was directed towards verifying her Cadence[®] simulations of an upcoming silicon cochlea chip, and suggesting improvements using the proposed modelling framework. In section 8.3, a selection of the results that we actually used for this process are presented in order to document some of the major findings, and to show the benefit of comparing chip simulations to numerical simulations in this manner¹. This chapter also serves as an example of how the details discussed in sections 4.4, 5.2, and 6.3 can be used to implement a nonlinear partition model in the framework.

The active VLSI models that Dr Hamilton proposed in her PhD (Hamilton 2008) to match the nonlinear response of the cochlea contained dynamics that allow a Hopf bifurcation. Eguíluz *et al.* (2000) suggested that the cochlea is poised at a Hopf bifurcation during its operation. Interestingly, cochlear models that incorporate a Hopf bifurcation share a close link to Van der Pol oscillator models, proposed in the 1980s. I noticed the connection when working with Dr Hamilton at UNSW in 2010, and contacted Hendrikus Duifhuis about his work on Van der Pol oscillator models to investigate the link further. He had also noticed the connection, and provided a proof of the relationship in his book (Duifhuis 2012), and an article (Duifhuis 2011). In section 8.1, we briefly overview Van der Pol oscillator models, Hopf bifurcation models, and analogue silicon cochleae.

Dr Hamilton sent the chip for fabrication in January 2011. A failure in one of the supporting circuits on the chip meant that we were not able to perform comparisons between the chip itself and results from the framework. This is particularly unfortunate, because Cadence simulations of the whole chip were not viable, so our comparisons before fabrication were limited to a small number of radial

¹For parties who are interested in determining the exact contribution to this combined work: Dr Hamilton's specialisation is in analogue VLSI circuitry, and my specialisation is in numerical simulations. All changes to, and results from, the Cadence simulations were Dr Hamilton's work, and all numerical simulations were my work.

segments. In place of comparisons between the chip and numerical simulations, we reproduce a figure from [Duifhuis *et al.* \(1986\)](#) in section 8.4 to show that mechanisms used to introduce the nonlinearities into the numerical model are correctly implemented. We select this particular figure because it shows an interesting similarity to physiological results in [Ren \(2002\)](#) for a cochlear model in which every segment is linearly unstable, and the nonlinearities are required for overall stability. Finally, we summarise the implications of this work for the proposed cochlear modelling framework.

8.1 Background

8.1.1 Van der Pol Oscillator Models

Strictly speaking, the Van der Pol oscillator takes the form:

$$\ddot{x}(t) - \varepsilon(1 - x^2(t))\dot{x}(t) + x(t) = F(t), \quad (8.1)$$

where $F(t)$ is a forcing function, and ε is tuning parameter. The term $-\varepsilon(1 - x^2(t))$ acts as a non-linear damping on the system. For the homogeneous system with $F(t) = 0$, the origin $(\dot{x}(t), x(t)) = (0, 0)$ is an unstable critical point if $\varepsilon > 0$ because $-\varepsilon(1 - x(t)^2) < 0$. The constant $-\varepsilon$ acts as a negative damping, adding energy into the system, but the nonlinear term $x^2(t)$ provides a level dependent damping that prevents the response from becoming unbounded. This balance between positive and negative damping drives the homogeneous system to a limit cycle around the origin.

In “Van der Pol” oscillator cochlear models, the nonlinear damping generally depends on $\dot{x}^2(t)$ rather than $x^2(t)$ resulting in the equation:

$$\ddot{x}(t) - \varepsilon(1 - \dot{x}^2(t))\dot{x}(t) + x(t) = F(t). \quad (8.2)$$

This is essentially the Rayleigh equation, as pointed out by [Duifhuis \(2011\)](#). That paper also describes the formal relationship between equations (8.1) and (8.2). Once again, the origin is unstable and the homogeneous system approaches a limit cycle. The fluid in the cochlear model couples these individual limit cycles so that each contributes to the forcing term of its neighbours (as well as its own, as appendix D shows). This prevents the segments from settling into a stable limit cycle, and the numerical results show a “chaotic” response (see [Duifhuis 2011](#)). This baseline response can be driven to the level of the background noise seen in the cochlea by manipulating the parameters of the segments.

While the idea of using Van der Pol oscillators as the model of radial segments was only suggested in 1980 (for example, see [Johannesma 1980](#)), models incorporating a nonlinear damping term predated these suggestions. For example, [Hall \(1974\)](#) studied the effect of both second- and third-order

damping nonlinearities on the generation of two-tone distortion products in a model. When that work was published, the first evidence of the nonlinear basilar membrane response was beginning to appear in the literature, but the evidence of active mechanisms in the cochlear partition had not been published. Therefore, the nonlinearities suggested did not consider an active response. (There was no negative damping included.) Van der Pol models on the other hand were proposed shortly after evidence of otoacoustic emissions that suggested the cochlea uses active mechanisms became available, which sparked the interest in the negative damping term (Duifhuis 2011).

8.1.2 Hopf Bifurcation Models

As mentioned in section 5.3, a (supercritical) Hopf bifurcation is characterised by a smooth transition in the response of the system from a stable equilibrium to a limit cycle. The transition is controlled by a set tuning parameters, and occurs at the “critical value” of these parameters (see Strogatz 1994). This implies that at least some of the tuning parameters in the set must be smoothly variable. When we study the dynamics of systems, we have the flexibility to treat even parameters that are fixed in practice as variable. (Considering the partition mass as a free variable, like we did in chapter 7, is an example.) However, the neural feedback from the brain to the cochlea suggests that there are tuning parameters available to enhance the response of the ear during the normal course of its operation. Eguíluz *et al.* (2000) took the suggestion that the brain tunes the response of the ear one step further: they suggested that the ear is maintained at a Hopf bifurcation point during its operation². Hudspeth *et al.* (2010) provides a recent review of the idea.

If a system is operated sufficiently close to a Hopf bifurcation, then certain key features of the response do not depend on the parameters of the system, but only the generic equation underlying the bifurcation. This fact has two important ramifications: in the presence of mismatch between components of a large system, independence of the parameters can be a significant advantage; and it allowed Eguíluz *et al.* to study a generic equation (normal form) for a Hopf bifurcation, but to comment on an entire class of cochlear models.

The generic equation they studied was:

$$\dot{z} = (\mu + i\omega_0)z - |z|^2z + Fe^{i\omega t}, \quad (8.3)$$

where $z(t)$ is complex, μ is a control parameter, ω_0 is the natural frequency of the homogeneous system, and $Fe^{i\omega t}$ is a periodic forcing function. The system undergoes a Hopf bifurcation for $\mu = 0$. Although equation (8.3) looks very different from (8.1) and (8.2), it is very closely related

²In section 3.3.4.1 we have noted that one of the central claims made in that paper, that the ear contains “essential nonlinearities” such that no audible sound invokes a linear response, may be in conflict with the best physiological data available. This may simply indicate that the system is not located perfectly at the bifurcation point but slightly below it, where a linear response can be obtained for sufficiently small inputs (see Eguíluz *et al.* (2000), page 5234).

(as shown formally in [Duifhuis 2011](#)). The equation:

$$\ddot{x}(t) + (\mu + \varepsilon \dot{x}^2(t))\dot{x}(t) + x(t) = F(t), \quad (8.4)$$

also has a Hopf bifurcation at $\mu = 0$ and can be converted to the generic form (8.3) if it is sufficiently close to the bifurcation point³. The similarity between this equation and expressions (8.1) and (8.2) is more apparent.

Therefore, three apparently different models share a close relationship. In the language of systems analysis, they study the same type of differential equation, but with different operating points: Hall's 1974 model operates in the region before bifurcation; the Hopf bifurcation models track the bifurcation point; and the Van der Pol oscillator models operate beyond the bifurcation. (The fact that cochlear models might need to operate in any of these three regions, and possibly cross between them, was a key motivator in the work following up on Matthews and Molnar's description of an instability in a numerical model near a Hopf bifurcation, described in section 5.3.1.1 and appendix D.)

8.1.3 Application to Silicon Cochleae

[Lyon and Mead \(1988a\)](#) proposed the first VLSI silicon cochlea. As motivation for their work, they cite the dual goal of using current understanding of the auditory system to build machines that hear (with replicating the cochlea as a first step), and using the relative success of the system to evaluate progress towards understanding hearing. They proposed an analogue cochlear model, implemented using CMOS transistors operated in the subthreshold region. (When transistors are operated in the subthreshold region they draw very little power, and the equations governing the device's operation have other useful features. The cost to operating in this region is that the components are sensitive to noise and mismatch.) The model that Lyon and Mead proposed drew on the Liouville-Green approach in order to obtain a 1D cascade of second order filters to be implemented on the chip. They also incorporated mechanisms for automatic gain control, implemented as variable negative damping, in the design.

Following the success of Lyon and Mead's 1988 VLSI silicon cochlea, a number of authors have extended the work in various directions. Dr Hamilton provides a thorough overview of the development in her PhD thesis ([Hamilton 2008](#)). Her own work extends work by [Watts \(1993\)](#) and [Fragnière \(1998\)](#), who introduced a 2D silicon cochlea model and a current domain 2D cochlea model respectively. She focused on matching the response of the nonlinear cochlea, and arrived at a silicon cochlea model that contained a Hopf bifurcation.

³However, [Hudspeth et al. \(2010\)](#) notes that this particular equation is "singular" in the sense that the imaginary part of its impedance is zero for all input frequencies at the bifurcation point. They argue that this means the model will not produce the "energy pumping" that might be critical to travelling wave amplification before the best place.

Silicon cochleae are an attractive method to implement and study the nonlinear models discussed in this section for a number of reasons. Numerical simulations of cochlear models are fundamentally deterministic, which means that it is hard to obtain the uncorrelated state necessary as a starting point for tests using the Van der Pol cochlear model in particular, but also for Hopf bifurcation cochlear models. The noise inherent in silicon chips ensures that the required uncorrelated state is easy to obtain. The mismatch between components that makes silicon cochlear models difficult to design is expected to be less of a problem for Hopf bifurcation cochlear models. Finally, nonlinear cochlear models are best studied in the time domain, which is computationally intensive to do in numerical simulations. Silicon cochleae run in real-time, allowing various inputs and parameters to be tested quickly (once designed and fabricated). These factors suggest that a silicon cochlear chip that can implement equation (8.4) on both sides of the bifurcation, with μ as a free parameter allowing for automatic gain control, will be a very useful device for studying these models. It will also be a significant step towards the “machines that hear” as envisioned by Lyon and Mead.

8.2 Implementation in the Proposed Framework

8.2.1 Select a Cochlear Partition Model Type

8.2.1.1 Model Inputs and Outputs

Once again, we need to decide whether to use pressure or potential as the solution variable. Dr Hamilton extends Fragnière’s 2D silicon cochleae that used pressure as the solution variable (see section 4.2.2 in [Fragnière 1998](#)). Therefore, we will also use pressure as the solution variable, which means that the input to the partition model is a pressure and the output is an acceleration. (Watts used potential as the solution variable in his silicon cochlea chip, so the potential based silicon cochleae are also viable, see [Watts 1993](#) section 4.1.1.)

8.2.1.2 The Dynamic Model

The partition model for this work is nonlinear and active. As it was argued in section 4.4.2, it is desirable for active mechanisms to build upon a passive model of the partition. The silicon cochlea chips use radial segments with second order dynamics⁴, so equation (4.33) is an ideal starting point:

$$\Delta p_{BM}(x,t) = M(x) \frac{d}{dt} v_{BM}(x,t) + R(x,t) v_{BM}(x,t) + K(x) \int_{t_0}^t v_{BM}(x,t) dt. \quad (8.5)$$

⁴Implementing fourth order dynamics would increase the footprint of each segment on the chip, reducing the number of segments that can be simulated.

In order to incorporate the nonlinear damping, we replace $R(x)$ by:

$$R(x,t) = R_p(x) - g(\mu - |v_{BM}(x,t)|^2), \quad (8.6)$$

where $R_p(x)$ is the original passive damping, $g \geq 0$ is the gain applied to the active mechanisms, and μ is a set-point to control the segments' proximity to the bifurcation point. With this expression for $R(x,t)$, the segment model has the same underlying form as equation (8.4). Bifurcation occurs when $R_p(x) = g\mu$. The specific choice of $R(x,t)$ in (8.6) matches the definition in [Hamilton et al. \(2010\)](#).

The dynamic model can be separated into linear and nonlinear parts:

$$a_{BM}(x,t) = a_{linear}(x,t) + a_{nonlinear}(x,t), \quad (8.7)$$

$$a_{linear}(x,t) = \frac{g\mu - R_p(x)}{M(x)}v_{BM}(x,t) - \frac{K(x)}{M(x)} \int_{t_0}^t v_{BM}(x,t) dt + \frac{\Delta p_{BM}(x,t)}{M(x)}, \quad (8.8)$$

$$a_{nonlinear}(x,t) = -\frac{g|v_{BM}(x,t)|^2}{M(x)}v_{BM}(x,t). \quad (8.9)$$

The linear part can be converted into a state update equation using the method shown in section 7.3.1.2. In this case, the state X_1 will be the segment velocity, and X_2 will be the segment displacement. To form the update equation of the nonlinear system, the contribution of the nonlinear part, which is $-gX_1^3/M(x)$, is added to the linear component of \dot{X}_1 .

The output linearisation equation (6.21) allowed:

$$a_{BM} = (\mathbf{C}_o \mathbf{X}_{dynamic} + D_o) + (\mathbf{E}_o \mathbf{X}_{dynamic} + F_o) \Delta p_{BM}, \quad (8.10)$$

$$= k_2 + k_1 \Delta p_{BM}. \quad (8.11)$$

a_{linear} determines both F_o in k_1 and \mathbf{C}_o in k_2 . These contributions are determined using the method shown in section 7.3.1.2. They are constant over time. We include the nonlinearity by defining D_o such that:

$$D_o = -gX_1^3/M(x). \quad (8.12)$$

D_o must be recomputed by a call to the `update_matrices` method each time k_2 is required.

8.2.1.3 The Data Model

In this case only a few data outputs were computed. They were: the velocity and displacement, X_1 and X_2 respectively; the feedback from the active mechanism, $g(\mu - |v_{BM}(x,t)|^2)v_{BM}(x,t)$; and the input pressure, Δp_{BM} . From equations (8.5) and (8.6) it can be seen that $g(\mu - |v_{BM}(x,t)|^2)v_{BM}(x,t)$ has the same units as Δp_{BM} , so comparing these values can indicate the significance of the active feedback in a given simulation.

8.2.1.4 Implementation

We implement this theory by extending the `OrganOfCortiComplexSegment` class to create a new class `HopfSegment`. The state update equation is implemented by replacing the virtual `dynamic_states_gradient_estimate` function in the `OrganOfCortiComplexSegment` class. The new method calls the linear method in the `OrganOfCortiComplexSegment` class to compute the linear component, and then added the nonlinear contribution.

We opted to provide two external inputs. These two inputs are linked with g and μ respectively, allowing these values to be changed even after a `HopfSegment` has been constructed – allowing an outer feedback loop to control the active mechanisms. The `update_matrices` method recomputes any values that depended on g or μ , as well as computing D_o for the output linearisation.

8.2.2 Select Tuning Parameters

The segment model described above requires five tuning parameters: $M(x)$; $R_p(x)$; $K(x)$; g ; and μ . The `Parameters` object does not need to supply a fluid density to the segment model in this case because the density occurs in the fluid model (see equation (4.15)).

This chapter investigates different mathematical models (with a common underlying structure) using the same segment class (`HopfSegment`). The different models required different parameter sets and numerous adjustments to the parameters were tested. For this reason, the parameters used for specific experiments will be reported as the tests are described.

8.2.3 Specify a Fluid Region

Where a fluid region was required for tests, it will be reported as the tests are described.

8.2.4 Choose Boundary Conditions

Section 4.5.1 discusses the importance of the stapes model, especially for nonlinear cochlear models. Including a stapes model allows the stapes-fluid boundary to respond nonlinearly. A model such as the one shown in equation (7.45) is sufficient. Although this chapter deals with nonlinear models, none of the models actually require a stapes model for correct operation, so it is omitted. (The mechanisms related to including stapes models using the framework were thoroughly tested in the “monolithic iterative” model described in appendix C, where published results were emulated.) Once again, the specific boundary conditions required for each test will be reported as it is described.

8.2.5 Assemble the Pieces

Once again, the modular approach encouraged by this framework facilitated construction of the models described in this chapter. The next section will focus on how small parts of the model were tested. The test described in section 8.4 adapted code originally used for the investigations described in appendix D. Only minor changes to that code were required. They involved replacing the parameters and segment objects constructed, adjusting tuning parameters, and formatting some output as desired for the new model. This illustrates the benefit of the modular approach that this framework aims to encourage.

8.3 Comparisons Between the Silicon Cochlea and the Modelling Framework

The object oriented approach adopted for this cochlear modelling framework encourages systematic testing of the components of a model as they are implemented and assembled. This procedure was also useful when testing the silicon cochlea model against the framework and visa versa. We performed the comparison at three levels: the cochlear fluid; the radial segment model; and small sets of segments interacting with the fluid model. The final test would be to compare results from the whole silicon cochlea chip against a numerical simulation, but the problem with one of the support circuits prevented testing at this level.

8.3.1 Cochlear Fluid

The Laplacian governing equation for the fluid model, $\nabla^2 p = 0$, only depends on the instantaneous boundary conditions. The expected fluid response for simple boundary conditions is easy to find from the governing equation. The modular structure of the framework allows the fluid response to be tested quickly when desired. Similar tests were conducted on the fluid model in the Cadence chip simulations, and they showed a frequency dependence in the fluid response.

In the silicon cochlea model, the fluid is implemented as a resistive grid that is equivalent to a finite difference mesh. The “resistors” that make up the resistive grid are actually implemented using pMOS transistors, because the resistivity of the silicon layers available in low cost processes is too low for micropower circuits (see [Watts 1993](#) page 81). The resulting resistive network is inherently frequency dependent. However, after noticing this problem, Dr Hamilton was able to adapt the parameters of the transistors to ensure correct operation of the network for the frequency range of interest in the cochlear model.

The findings of section 7.2 also have implications for the silicon cochlea: if the ratio of fluid density to segment mass is incorrect, the interaction of the two models will be misrepresented. Fragnière discusses this problem in section 4.4.6 of Fragnière (1998). His figure 4.12 shows good agreement with our results in chapter 7.

8.3.2 Cochlear Partition Model

8.3.2.1 Desired and Undesired Nonlinearities

When studying a circuit that should be linear, any nonlinearities that appear in the response are clear evidence of a problem. Transistors that are operating on the boundary of the subthreshold regime or signals that hit one of the supply rails both introduce undesired nonlinearities into the response. If the circuit is known to be nonlinear, then a nonlinear response does not indicate a problem in and of itself. Numerical simulations of the mathematical model behind the circuit can provide a useful means to establish the type of response expected from the circuit. Comparison between the numerical simulation and the circuit implementation quickly show whether a specific form of nonlinearity in the response is desired or undesired.

Figure 8.1 shows results that Dr Hamilton produced from her Cadence simulations. Panel (a) shows how the segment output evolves over time for a small input. The output grows steadily until it reaches a limit cycle after approximately 20 ms. The settled velocity waveform, panel (b), does not show any significant nonlinearities to the eye. For a large input, panel (c), there are obvious nonlinearities in the response. The question is, which nonlinearities are desired, and which are not?

In order to investigate this question, a single `HopfSegment` was simulated. The parameters selected were: $M(x) = 0.15 \text{ g} \cdot \text{cm}^{-2}$; $R(x) = 200 \text{ dyn} \cdot \text{s} \cdot \text{cm}^{-3}$; and $K(x) = 5.4 \times 10^7 \text{ dyn} \cdot \text{s}^2 \cdot \text{cm}^{-3}$. These values correspond to the point $x = 1.46 \text{ cm}$ in the data set from Neely (1978), which was also used in Fragnière (1998). The characteristic frequency of this segment (in isolation) is 3017 Hz, so this frequency was used for the input. The values of g (1×10^6) and μ (300×10^{-6}) were adjusted to match the response in panel (a) of figure 8.1. The response of the numerical simulation using these values is shown in figure 8.2 on page 186. In panels (a) and (b) the input amplitude was $0.05 \text{ dyn} \cdot \text{cm}^{-2}$, and in panel (c) it was $20 \text{ dyn} \cdot \text{cm}^{-2}$. The results show good overall agreement with the Cadence simulation. However, when panel (c) of the two figures are compared, it can be seen that the rising edge of the response has qualitative differences, while the falling edge does not. This shows that some undesired nonlinearities are affecting the rising edge of the waveform in the Cadence simulation. Having identified the undesired nonlinearities, they could then be investigated further.

It is important to note that the cochlear partition governing equations (8.5) and (8.6) show some harsh nonlinearities for large inputs or set-point $g\mu$. Figure 8.3 on page 187 shows the acceleration




Figure 8.1: Cadence simulation of a single segment for a small input, panels (a) and (b); and a large input, panel (c). In panel (c), note the difference between the shape of the falling edge and rising edge of the wave.

of the segment for the two inputs used above in (a) and (b) respectively, and $1000 \text{ dyn} \cdot \text{cm}^{-2}$ in (c). Panel (a) shows no visible nonlinearity. (b) and (c) show obvious, and even extreme, nonlinearities. Given that distortion products are as much as 40 dB smaller than the input signal, it seems unlikely that the cochlear partition acceleration is ever as nonlinear as this partition model allows. More recent cochlear partition models often include a tanh characteristic that would prevent these extreme nonlinearities (see [Liu and Neely 2010](#) for example).

The harsh nonlinearity can have an effect on the time stepping method used in numerical results. The rapid transition in the state around the response maximums and minimums require a small time step. Although silicon cochlea chips are difficult to design due to the limitations on signal sizes and component values, one major advantage they have is that they are inherently an analogue solution. As such, they neatly sidestep any issues related to time stepping methods and stiffness of the model.

8.3.2.2 Implementation of Gain

In a silicon cochlea chip, any circuitry that is added to the radial segment design increases its foot print, and ultimately reduces the number of segments that can be included in the model. In order to reduce the circuitry required, the gain g was generally set to a fixed value in previous active cochlea chips. The numerical simulations with the framework highlighted the importance of the




Figure 8.2: Numerical simulations of a single segment for sine wave input at 3017 Hz, the best frequency of the segment. The input amplitude for panels (a) and (b) is $0.05 \text{ dyn}\cdot\text{cm}^{-2}$, and for panel (c) it is $20 \text{ dyn}\cdot\text{cm}^{-2}$. These results show good agreement with the silicon cochlea results in figure 8.1, except for the rising edge of the wave in panel (c). This indicates a problem in the silicon cochlea that can be investigated further.




Figure 8.3: Numerical simulations of a single segment for sine wave inputs at 3017 Hz, the best frequency for the segment. The input amplitude for panels (a), (b), and (c) are $0.05 \text{ dyn} \cdot \text{cm}^{-2}$, $20 \text{ dyn} \cdot \text{cm}^{-2}$, and $1000 \text{ dyn} \cdot \text{cm}^{-2}$ respectively.

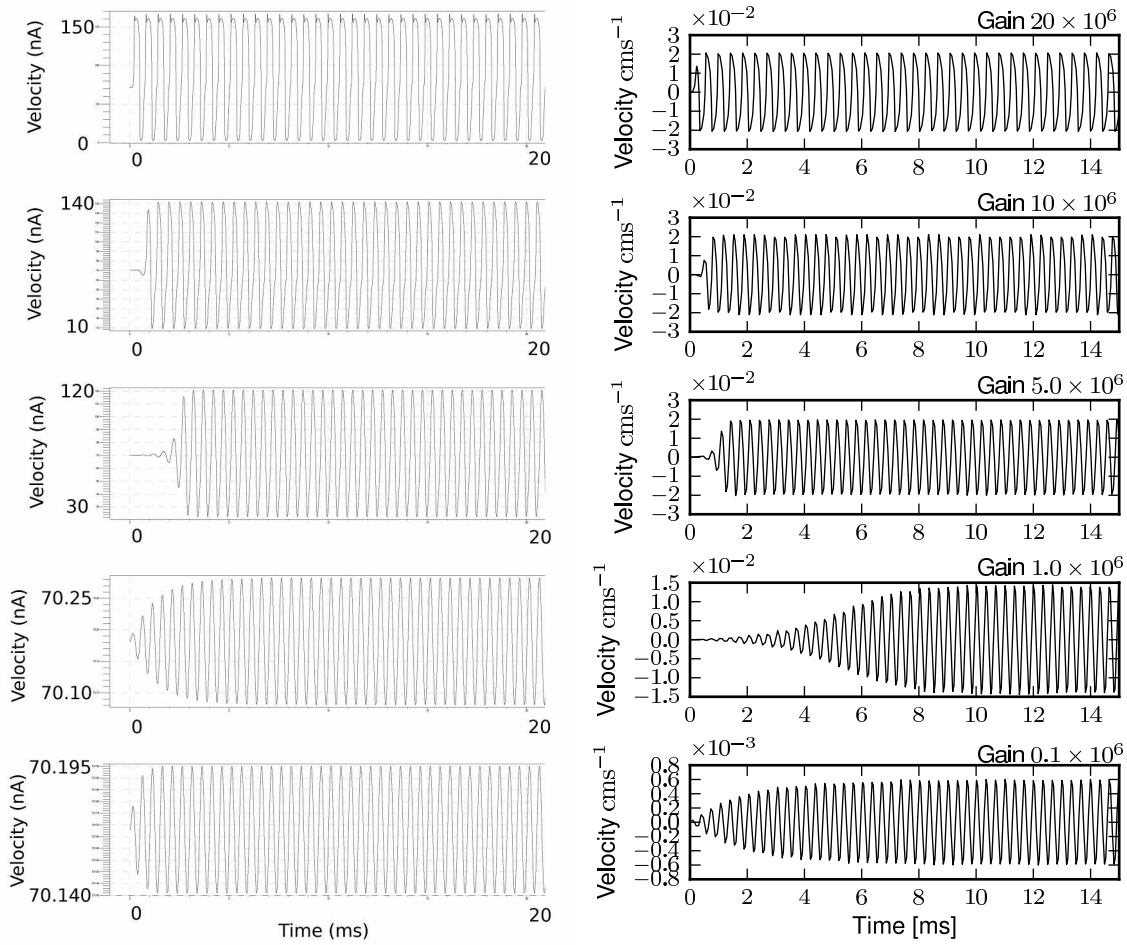


Figure 8.4: Comparison of the effects of the gain parameter on the Cadence simulation (left) and framework implementation (right). The qualitative agreement found confirms that the gain has been correctly implemented. Note that the gain g interacts with μ to determine the proximity of the segment to the bifurcation point. This results in the change in the response shape with gain. The apparent increase in limit cycle size with increasing gain is also an artifact of their interaction. For fixed $g\mu$ increasing g decreases the limit cycle size.

gain for controlling the size of the limit cycles that occur if a radial segment is operating above the bifurcation point. This motivated implementing the gain as an extra control variable in the silicon cochlea. Figure 8.4 shows the response of the radial segment model for various gains. The left-hand panel show Cadence simulations, and the right-hand panel shows numerical results. $M(x)$, $R(x)$, and $K(x)$ were as described in section 8.3.2.1. In this case, $\mu = 350 \times 10^{-6}$ was used. Gain values in the range 1.0×10^5 to 20×10^6 provided a qualitative match to gains of 1 nA to 5 nA in the Cadence simulation. The close similarity in the form of the response shows that the gain has been correctly implemented in the chip, and it is having the desired effect. Note that changing the gain also has an effect on the proximity of the system to the bifurcation point, determined by $g\mu$. Therefore, increasing the gain appears to increase the size of the limit cycles in the figure. If μ is adjusted to keep $g\mu$ constant as g is increased, the size of the limit cycle decreases, and the shape of the response




Figure 8.5: A comparison between two-tone suppression in an active silicon cochlea and physiological results showing linearisation in the response to high stimulus intensities. (From [Hamilton 2008](#), figure 7.28. Right-hand panel originally appeared as figure 3 in [Ruggero et al. 1992](#).)

also remains the same, instead of varying as seen in the figure.

8.3.2.3 Linearisation of the Cochlear Input/Output Curves

It has been suggested that at high stimulus intensities the stapes input supplies sufficient energy to evoke a detectable response, and the nonlinear active mechanisms are no longer required. Furthermore, the amount of energy that the active mechanisms can produce is presumably limited, and becomes negligible at these intensities, so the response of the cochlea should become linear.

[Robles and Ruggero \(2001\)](#) review this idea (see section II.A.1 in that paper). They list some physiological evidence that supports it by showing that the compressive input/output function slope at low stimulus intensities approaches linearity around 90–100 dB SPL. However, they also list some studies that found a compressive response even as high as 100 dB SPL, and note that the linearisation observed in other studies might be due to cochlear damage that can occur due to exposure to high sound levels.

With this note of caution aside, we observe that the dynamic model of the partition used in this chapter remains compressive for high stimulus levels. [Eguíluz et al. \(2000\)](#) equation (1) gives the solution to equation (8.3) for the 1:1 locked scenario (that is when the input and output frequencies are the same). For large inputs, the response varies as $F^{\frac{1}{3}}$, showing a compressive characteristic. Numerical results from the `HopfSegment` class confirmed the compressive response.

On the other hand, Dr Hamilton has observed linearisation of the cochlear response in previous silicon cochlea chips (see figure 8.5). This presumably occurs because some internal signal either hits one of the supply rails or moves outside the sub-threshold regime, limiting the compressive response. These potential causes share interesting similarities with the original suggestion of why linearisation of the response might occur in the cochlea. If, however, the linearisation of the response




Figure 8.6: Cadence simulation of four radial segments coupled by fluid. Each panel presents the results for one segment, with characteristic frequency (CF) indicated. Four inputs to the model were simulated, each corresponded to one of the characteristic frequencies.

that has been observed is only associated with cochlear damaged, then any transition to a linear response in silicon cochlea chips should occur outside the input range of interest.

8.3.3 Four Coupled Segments

Figure 8.6 shows results from a Cadence simulation test with four radial segments, coupled by fluid. The characteristic frequencies of the four segments were 420 Hz, 780 Hz, 1500 Hz, and 2900 Hz. The final stage of the fluid model was grounded representing a zero pressure condition there. Dr Hamilton reported that the simulation struggled to find a d.c. operating point, which restricted the range of parameters that she could test. Can the numerical model be used to comment on the results obtained?

The velocity response in figure 8.6 does not show any visible nonlinearity, and the rapid settling time of the response suggests that the segments have a low Q . ($Q = 1.0$ was ultimately used.) Both of these factors suggest that a linear model should be able to reproduce these results. (That is, $g = 0$.) Segments with comparable characteristic frequencies were obtained by defining an exponential




Figure 8.7: Numerical simulation of four radial segments coupled by fluid. Each panel presents the results for one segment, with characteristic frequency (CF) indicated. Four inputs to the model were simulated, each corresponded to one of the characteristic frequencies.

cochlear mapping parametrised by the length of the section, L :

$$f(x) = 4000e^{-2.58x/L}. \quad (8.13)$$

The wide range spanned by the characteristic frequencies would suggest that L should be a large fraction of the actual cochlear length. However, this resulted in undesired low frequency whole cochlear resonances, probably associated with the apical boundary condition. For a shorter section length $L = 0.5$ cm, whole cochlear resonances did not occur. This value is comparable to the length that 5 fluid divisions in a silicon cochlea would model. Equations (7.1) and (7.43) were used to determine the stiffness and damping from a value for M . ($M = 1.0 \text{ g} \cdot \text{cm}^{-2}$ was ultimately used.)

Large variations on all of these parameters were tested by trial and error to determine their effect on the simulation. The reported parameters appeared to provide the closest match to the main visible features. Figure 8.7 shows the match attained. The high frequency response around 0 s was never seen in the numerical simulations, and it is presumably either associated with the dynamics of the circuit components, or an artifact due to the simulator struggling to find a d.c. operating point. The mass value that provided the best match to the circuit behaviour is of concern, because figure 7.1

would place it close to the the high mass scenario described in section 7.2.3. However, given that the overall match is relatively poor, and that M and L are likely to interact, further tests are required to confirm whether it is in fact a problem.

This experiment highlights another important way that numerical simulations can support the development of silicon cochlea chips: they can be used to determine suitable conversions between the units in the chip to the units in the literature for a given tuning.

8.4 Operating Cochlear Models Beyond the Bifurcation Point

In place of comparisons between the full silicon cochlea model and the numerical implementation using this framework, we show some results for the Van der Pol oscillator model described in [Duifhuis *et al.* \(1986\)](#) that were reproduced using the framework. By replicating published results, we demonstrate that the mechanisms supporting nonlinear models in the framework have been correctly implemented. (This has also been demonstrated in appendix D, which studied the behaviour of a model with four states in the vicinity of the bifurcation point.)

In this implementation, the MKS unit system was adopted. The cochlear fluid was rectangular with length 0.035 m and height⁵ 0.001 m. The density is then $1000 \text{ kg} \cdot \text{m}^{-3}$. The finite element grid consisted of 200×5 finite element cells (giving 400 radial segments on the cochlear partition). A time step of 2.5×10^{-6} s was used. $R(x)$ was specified by:

$$R(x, t) = -d_1 + d_2 v_{BM}^2(x, t), \quad (8.14)$$

where $d_1 = \varepsilon \sqrt{M(x)K(x)}$, $d_2 = d_1 \times 10^{12}$, and $\varepsilon = 0.05$ – rather than the form shown in equation (8.6). These values can be related back to the parameters required for the HopfSegment using the expressions: $R_p(x) = 0$; $g = d_2$; and $\mu = d_1/d_2$. The original model was 1D and our implementation is 2D. Appendix C shows that tuning differences between 1D and 2D models are to be expected. The original mass was used ($M(x) = 0.5 \text{ kg} \cdot \text{m}^{-2}$, constant along the partition length) and the stiffness was adjusted to account for the dimension. The stiffness was given by:

$$K(x) = 10 \times 10^9 e^{-300x} \text{ kg} \cdot \text{m}^{-2} \cdot \text{s}^{-2}. \quad (8.15)$$

A zero pressure boundary condition was imposed at $x = 0.035$ m. At the stapes, a sine wave input velocity of $7.5 \times 10^{-3} \text{ m} \cdot \text{s}^{-1}$ at 1 kHz, with a 2 ms raised cosine onset, provided a good match to the 30 dB input described in [Duifhuis *et al.* \(1986\)](#). Figure 8.8 shows the match attained. [Duifhuis *et al.* \(1986\)](#) discuss the importance of using a middle ear model for nonlinear cochlear models in general, but argue that it is not required for the 30 dB input due to the small size of the response evoked. The

⁵This height value did not appear in the original paper, and was supplied based on the basilar membrane length.




Figure 8.8: Results from the Van der Pol oscillator model described in [Duifhuis *et al.* \(1986\)](#). Parameters of the original 1D model were adjusted for the 2D implementation in the framework. Compare the restricted range of the response to [Ren \(2002\)](#), figures 1c and 3c.

initial conditions were as for [Duifhuis *et al.* \(1986\)](#): X_1 (the velocity) alternated between +1 and -1 nm/ms along the partition.

[Duifhuis *et al.*](#) comment on the fact that adjacent segments can have opposite phase in these models. They note that because the hair cells are spaced by $\approx 10 \mu\text{m}$, and the deflections are on the order of nanometres, opposite phases for adjacent segments are still physically plausible. The finite element method uses shape functions that assume continuity in-between values at nodes. This assumption would be violated by these models in the limit of an infinitely refined mesh but, as we have already noted in section 4.4.3, there are other good reasons not to push the models to this extreme. Therefore, using the finite element method for these models is internally consistent for reasonable segment sizes.

This specific example was chosen as the test for models operating above the bifurcation point because it also allows us to return to discuss the criticism against travelling wave theory that was covered in section 7.1.3.3. [Ren \(2002\)](#) found evidence for a travelling wave on the basilar membrane over a restricted area ($\approx 600 \mu\text{m}$ for a low level 16 kHz tone). In a Van der Pol oscillator cochlear model, every point on the basilar membrane is slightly unstable resulting in limit cycles at all segments, with little long term correlation. Figure 8.8 shows that the limit cycle can be driven to the level of background noise, in which case an input signal causes the otherwise uncorrelated response to be entrained over a small region. This allows a travelling wave to first emerge some distance from the stapes. Therefore, the 1986 model by [Duifhuis *et al.*](#) suggests a possible explanation for [Ren's](#) recent findings⁶.

8.5 Summary

8.5.1 VLSI Implementation and Hopf Bifurcation Models

In this chapter, we have investigated Hopf bifurcation and Van der Pol oscillator cochlear models using the framework. The main goal of these investigations was to support the development of a VLSI silicon cochlea. We found that numerical simulations can support the design process in a number of ways. They can be used to provide examples of the desired response of the mathematical model under various conditions, confirming where the VLSI implementation is correct and highlighting deviations from the expected behaviour. They can also facilitate in the comparison between results from the VLSI implementation and the literature by determining equivalent tuning points between the implementation and the model, providing the required unit conversions in the process.

The investigations were largely successful. They suggested a problem with the fluid model in the VLSI implementation that could be rectified by a simple retuning of the circuit. The numerical re-

⁶A reviewer notes that this finding might also be the result of a limited signal to noise ratio in the experiment (see [Ren \(2002\)](#), pages 17102–17103 for a discussion of the noise floor in the experiment).

sults helped to distinguish between desired and undesired nonlinearities in the circuit, highlighting problems to be addressed. They also confirmed that the “gain” control provided in the mathematical model was correctly implemented. Tests using four coupled segments provided mixed results. They suggested that the segments were coupled in a meaningful way and were providing sensible output, but suggested that the fluid density to segment mass ratio that is implicitly set by the circuit parameters might need to be adjusted. The difficulty Dr Hamilton had obtaining Cadence simulation results for this test prevented a firmer conclusion to be drawn on this point.

The results presented predict that the current VLSI implementation should be able to accurately operate at, or on either side of, the bifurcation point. We are therefore hopeful that this chip will be able to contribute to our understanding of the behaviour of these models in the face of noise. Numerical simulation, in particular of Van der Pol oscillator models, are restricted by the deterministic nature of computations. Furthermore, numerical models do not run in real-time. Circuits inherently contain background noise and do run in real-time. This will facilitate rapid testing of the model once the chip becomes available. A failure in one of the supporting circuits in the chip is currently delaying further testing.

8.5.2 Van der Pol Oscillator Models

In place of direct comparisons between the chip and numerical simulations, we reproduced certain published results from a Van der Pol oscillator model to demonstrate that the framework can implement models that operate beyond the bifurcation point correctly. The results from this model show intriguing similarities to recently published physiological results, justifying our continued interest in these types of models. A recent publication discussed the close relationship between Hopf bifurcation models and Van der Pol oscillator models. We would go further to claim that, under some circumstances, they can actually be a single mathematical model, with different choices of operating points. Our results show that the same mathematical model can operate in various regimes, supporting this claim.

8.5.3 Support for Central Ideas in this Thesis

This work also shows the value of the central premises of this thesis. The paradigm of mathematical versus implemented models was useful in this comparison between a VLSI implementation and a numerical implementation of the same mathematical model. Furthermore, the fact that we dissociate models from their implementations, and hence tuning parameters, makes obvious the conclusion that Hopf bifurcation models and Van der Pol oscillator models can be simply different operating points of the same mathematical model.

This work again showed that comparisons between different implementations of the same mathematical model can be tedious. This provides motivation for a framework that is fair to a wide range of models, and allows them to be compared directly. However, this chapter also balances that point by demonstrating the benefits that can be reaped by comparison between different implementations. VLSI implementations have strengths at the precise point where numerical implementations are weak: time stepping methods. By comparing implementations and trying to justify any differences that are seen, we are performing important verification of the model. In fact, verification itself can be viewed as the process of comparing our new implementation against the “perfect implementation” provided by the mathematical model.

In order to obtain a modelling framework that is fair to a wide range of models – including some that have not been proposed yet – flexibility and extensibility is essential. In the proposed framework this is achieved through the use of modular, object orientated programming. Modularity is also a central topic in electrical engineering, so it is not surprising that comparisons between small parts of the VLSI and numerical implementations were feasible and useful. The object orientated approach taken in this framework proved its worth by facilitating the diverse investigations that are reported in this chapter.

Chapter 9

Conclusions

This concluding chapter is divided into three sections. The first “track” running through the thesis, which motivates and describes an extensible cochlear modelling framework, is discussed in section 9.1. Section 9.2 serves as an interlude, where the various models that have been implemented in the framework are discussed. Potential future directions are also noted. Finally, the results that will have general interest (our second track) are summarised.

9.1 An Extensible Cochlear Modelling Framework

In this thesis, an extensible cochlear modelling framework has been proposed. The framework relies on the flexibility inherent in the finite element method and state space models to ensure that a wide class of cochlear models can be implemented. Adopting the principles of object orientated programming provides the required extensibility.

Reasoning about the process of modelling itself suggests that comparisons between models are required, and will continue to be needed in the future. An analogy of the process of cochlear modelling using three triangles has been presented. The first is the mathematical idealisation of a triangle and the other two are separate approximations to that ideal. It can be used in a variety of ways to understand the process. Sometimes our objective is to reproduce the biological cochlear as accurately as possible, and at other times we seek simplicity in our models, reminiscent of the mathematical triangle. Ultimately, a single cochlear model will not satisfy all the diverse uses to which it could be put. A framework that can be used to compare competing cochlear models will be useful now and into the future. The paradigm of separating a model, that is its mathematical abstraction, from its implementation has proven to be a useful way of thinking about the variety of models of the cochlea, and their implementation.

The scope of the thesis narrowed steadily from these broad, fairly philosophical considerations to focus on a specific set of cochlear models – lumped parameter models that use a Laplacian fluid

formulation. This allowed the whole process of modelling to be discussed, and firm examples of implementations to be given. Settling on a specific model and implementation requires a number of trade-offs to be balanced. The trade-offs have been discussed, and opportunities to relax restrictions on the modelling framework have been noted. The physiology of the cochlea allows us to impose certain restrictions without concern. Some central design choices, such as using a time domain modelling approach without specifying a limited set of allowable nonlinearities up front, place unavoidable restrictions on the model that must simply be accepted. The characteristics of the fluid-structure interaction in the cochlea places undesired restrictions on the independence of classes in the framework. However, these same characteristics also imply that the style of models that we have investigated are representative of a wide class of cochlear models – guaranteeing that the extensibility of this framework can be achieved.

Throughout the course of the thesis, the value of an extensible framework that allows different models to be compared directly has been seen. Slight differences that arise from implementing even the *same* model using different approaches in different technologies can confuse the issue when the performance of a cochlear model is assessed. This is not to say that it is bad to implement models in various ways. On the contrary, comparisons between different implementations for the same model provided important insights and verification of the framework. It is simply a principle of good scientific investigation to allow only the minimum number of factors to vary between experiments. An extensible cochlear modelling framework offers this type of control. The object orientated approach adopted in the framework allowed tested code to be quickly adapted to implement new cochlear models. The methodical process of testing code class by class that is encouraged by object orientated programming proved to be a natural parallel to the modular approach that is used in electrical engineering. This facilitated the comparisons between the numerical simulations and VLSI cochlear chips that were performed.

9.2 Models Implemented in the Framework To Date

In the course of this work, a variety of cochlear models have been implemented in the time domain using the proposed framework. Firstly, a linear second order (active) model typical of classical cochlear models was used to demonstrate the process of assembling models in the framework, and to comment on the question of whether there is a true, or only an apparent, travelling wave in the cochlea. In the process, an analysis of the segment to segment coupling through the cochlear fluid was developed. The mechanisms used to swap a simulation from a single scala model to a model with two separate scalae were also tested with this model.

A nonlinear second order cochlear model that can be used to study either the Van der Pol oscillator or Hopf bifurcation cochlear models was constructed. This model also served as an example of how nonlinear cochlear models can be implemented in the framework. It was compared to Cadence[®]

simulations of a VLSI silicon cochlea chip that implements the same model, and the results were used to verify and improve that chip.

A linear fourth order active cochlear model was implemented during tests to verify that the modelling framework could correctly implement the monolithic model. Results from this model were favourable: not only did the framework reproduce the monolithic model to numerical accuracy, but the work showed that the “monolithic iterative” style of computation used in the framework is an efficient way to solve cochlear models.

Finally, a nonlinear fourth order active cochlear model was implemented to investigate a description of instability in the literature. The instability turned out to be of no concern for the modelling framework, but it provides an interesting example of how the slight differences between two implementations of the same model can confuse comparison. This model was investigated further with a first attempt at partitioned approach, designed under strict predetermined limitations.

All of these models have been implemented in the framework using a 2D fluid model, although comparisons to 1D models were described at one point in the thesis. Provision for feedforward mechanisms have been made in the framework, but they are not fully implemented. Their implementation will not be finalised until a model requires them, to avoid overdesigning these mechanisms. Building a model of this type would be a natural next step in the development of this framework. On the other hand, the mechanisms for external inputs to and from the cochlear model have been implemented and used. These mechanisms can be used to “close the loop” and place, say, a Hopf bifurcation cochlear model under feedback control. Equally, given that two scalae cochlear models can be constructed in the framework, the all the major requirements for Matthew Flax’s active compression wave cochlear amplifier seem to be met.

Overall, the modelling framework has been used to implement a range of models. At this point, various directions of further development are available.

9.3 Supporting Derivations and Results

Three results that were derived while developing this framework seem to be of general interest to the modelling community. Firstly, it has been shown that the state space modelling approach introduced by Elliott *et al.* is related to the previous modelling approach proposed by Allen and Sondhi by an identity for the inverse of the sum of matrices. This result allowed direct comparisons between the framework and two styles of monolithic computations, allowing important verification of the framework to be performed. The result will be of wider interest, because it allows researchers to select the more convenient of two starting points when deriving analytical results.

The second key result is that the linearisation of a nonlinear model shares certain key features to a linear model. (In particular, the form of the fluid coupling matrix is unchanged.) This is an excel-

lent result for the framework, because it shows that there is an inherent consistency between linear models, in which the fluid-structure coupling has been analysed in detail, and nonlinear models. The result will also be of general interest, because it extends the monolithic cochlear models, mentioned above, to nonlinear models. A given nonlinearity will affect some or all of the terms in the linearisation. The specific terms affected determine the type of fluid solver required for a model. Therefore, using the linearisation, the effect that various classes of nonlinearities will have on the suitability of a given modelling approach can be assessed.

Finally, a compact proof that the cochlear fluid can provide local segment to segment coupling has been provided. The proof extends one of the expressions for the monolithic model described above. Despite strong evidence from numerical cochlear models and analytical techniques like the Liouville-Green approach, the logic that because the cochlear fluid is incompressible, all points in the cochlea must receive the same input signal, continues to be used. Ironically, the incompressibility *does* couple all points on the cochlea, but when it interacts with a deformable boundary with similar density to water, then the closed loop behaviour of the coupling along the boundary is similar to the coupling within the fluid itself – it is strongly dominated by local effects. I do not think that this local coupling automatically rules out models that provide alternatives to the travelling wave hypothesis, but it does show the importance of using macroscopic cochlear models when testing theories of hearing.

Therefore, this cochlear modelling framework has been used to implement and test a variety of cochlear models. There are a range of directions that are available for future development of the framework. Furthermore, it has also been used to prove results of general interest. I trust that the framework itself, and the results that I have shown while using it, will assist in our quest to understand the intricacies of the cochlea and our sense of hearing in general. I look forward to seeing the medical treatments and biologically inspired technology that will result from continued work in this field.

Appendix A

Modes of OAE Observation

David Kemp, who first observed otoacoustic emissions from the cochlea, provides a summary entitled “Modes of OAE observation” in the appendix of a review article [Kemp \(2008\)](#). Since this provides such a clear description of the terms and evidence of their cochlear origin, it is reprinted below:–

The commonly used modes of OAE observation are listed below, each denoted by its usual acronym and followed by a summary of:

- the nature of the stimulus applied;
- the usual method of analysis used to separate it from other signals with any additional procedures needed;
- the nature of the OAE signal;
- the means of confirming its cochlear origin.

SOAE: Spontaneous otoacoustic emission. No stimulus is applied. Narrowband frequency analysis is used for signal recovery. Activity is identified statistically by its level above the background noise. SOAEs in humans usual consist of a stable tone with a bandwidth (frequency stability) of less than 1 Hz. A cochlear origin is confirmable by frequency-selective suppression of the emission by an applied tone, which reveals a sharp tuning curve.

SFOAE: Stimulus frequency otoacoustic emission. A single-tone stimulus is applied. Narrow-band frequency analysis is used to remove noise. Separation of the SFOAE from the stimulus is on the basis of the OAE nonlinear growth properties and is achieved by scaled vector subtraction linked to stimulus-level manipulations or, alternatively, the introduction of a second suppressor tone of a different frequency. SFOAEs consist of a pure tone at the stimulus frequency. A cochlear origin is confirmable by frequency-selective suppression by an applied tone, which reveals a sharp tuning curve.

TEOAE, CEOAE: Transient or click-evoked otoacoustic emission. A short-duration stimulus (<3 ms) is applied repetitively, either a wideband or frequency-limited stimulus. Synchronous averaging is used for signal recovery from noise and is followed by frequency analysis. Separation from the stimulus and the ear-canal-plus-middle-ear acoustic responses is usually by time gating and/or waveform subtractions linked to stimulus-level manipulations exploiting the nonlinear growth of OAEs. TEOAEs consist of a complex frequency-dispersed oscillation consisting mainly of the frequencies present in the stimulus. True TEOAE exhibit both time- and frequency-dependent suppression in response to additional stimuli [Kemp and Chum \(1980\)](#).

DPOAE: Distortion product otoacoustic emission. Two stimulus tones are applied, normally less than a half octave apart. The response is identified statistically by level above the background noise at neighboring frequencies. DP frequency is precisely related to the stimulus frequencies f_1 and f_2 by the formulas $f_1 - N(f_2 - f_1)$ for the lower band and $f_2 + N(f_2 - f_1)$ for the upper side band (N an integer greater than 1). DPOAE are often thought of as pure-tone emissions, but this is only true when a single component is separated by processing. Composite DPOAEs (taking all components) are complex modulated sine wave oscillations pulsing at the stimulus difference frequency (see [Kemp 2008](#), figure 1.3 and [Kemp et al. 1986](#)). Similar distortion-product generation is common in electromechanical systems. A cochlear origin is confirmable by frequency-selective suppression by an applied tone, which reveals a tuning curve, and also by determination of the delay.

Useful though this classification is for practical purposes as evidenced by its presence in most texts on the subject of the concept of OAEs, the use of the term “types” distracts from the subject of this chapter, which concerns the origins of OAEs in the cochlear process.

End of quotation from [Kemp \(2008\)](#).

Appendix B

Results from Parallel Processing Tests

B.1 Introduction

In this section, we discuss the results from tests to determine the potential speed up that parallelisation can provide for the cochlear models that were studied in this thesis. The code described in chapter 7 was used in these tests.

The preliminary tests, described below, show that the fluid solver dominates the run-time of the program, so some background on the types of solvers that can be used is appropriate. There are two basic approaches to solving large linear systems: direct methods such as Gaussian elimination (LU factorisation is an efficient method to perform Gaussian elimination); and iterative methods such as Krylov space techniques, and the Conjugate Gradient method in particular. Sparse LU factorisation is known to be an efficient method of solving the linear systems that arise from 2D finite element codes if the matrix equation consists of less than about 100 000 unknowns (Strang 2007, chapter 7). deal.II provides wrapper functions for UMFPack, which provides this type of solver (Davis 2004). UMFPack is essentially a serial solver, unless the Basic Linear Algebra Subroutines (BLAS) implementation provided on a system is parallelised. At the present time, parallel BLAS implementations are not common. An LU factorisation can be used multiple times to solve a system, provided that the system matrix (\mathbf{K}_{aug} from equation (6.15) in our case) remains unchanged. In the particular example that we study, \mathbf{K}_{aug} does not change, but this property does not hold in general.

Preconditioned Conjugate Gradient (CG) solvers are the Krylov space solver of choice for symmetrical positive definite matrices (Strang 2007, chapter 7). \mathbf{K}_{aug} meets these criteria. CG solvers perform a series of iterations until an error bound on the solution is satisfied, but LU factorisation obtains the solution to numerical precision directly. Specifying a smaller error bound for the CG solver reduces the final error in the solution, but requires further iterations increasing the time taken to solve the system. Therefore, whether a CG solver outperforms a solver using LU factorisation is

affected by the size of errors that can be tolerated in the solution. A good preconditioner improves the rate of convergence of iterative solvers, thereby reducing the number of iterations required to reach the error bound and hence the run-time. However, there is a trade-off in that constructing and applying the preconditioner takes time.

Incomplete LU factorisation can be used as a preconditioner for CG. (We opted to use other preconditioning options, because full LU factorisation was already considered.) deal.II provides a range of preconditioners, including the following: Jacobi; Chebyshev; and symmetric successive over-relaxation (see the deal.II documentation, [Bangerth and Kanschat 2011](#)). For the system studied, the symmetric successive over-relaxation (SSOR) preconditioner with the relaxation parameter set to 1.8 performed the best of these three preconditioners. However, the deal.II implementation of this preconditioner is not parallelised. Matrix multiplications by a vector in deal.II are parallelised, and the unpreconditioned CG solver in deal.II benefits from this parallelisation.

Trilinos provides an MPI based CG solver and various preconditioner options. One of those options is SSOR, and once again 1.8 proved to be the best relaxation parameter setting. The Trilinos preconditioner is parallelised, but applies the preconditioner blockwise, rather than to the matrix as a whole. While this makes the preconditioner easier to parallelise, it also reduces the performance. When working in a message passing environment, the data and solution are distributed over all the cores. Collecting this data back to a single core (say to write to file) is expensive, because it requires communication between all the cores. Well parallelised code tries to avoid this where possible, but it is sometimes unavoidable. To obtain linear speedup, the run-time must halve each time the number of cores doubles. Superlinear speedup (where the run-time reduces faster than linear) can occur on some problems, but it tends to indicate that the test problem is too small, and fits in the processor cache.

By studying the performance of these various solvers in one of the codes used for this project, we will see that while parallelisation is an important topic for cochlear modelling in general, it is currently unnecessary for the types of models studied in this thesis.

B.2 Method

We first assessed the performance of the code before parallelisation. To do this, all nonessential methods (such as the postprocessing of the solution described in chapter 7) were disabled. The remaining code constituted the minimal number of instructions required to initialise and solve the system. We verified that the solution was unchanged by this process. Using the deal.II `TimerOutput` class, we measured the time that the program spent performing three basic types of operations: assembling the matrix equation at each time step; processing calls to the framework; and solving the system. We did not explicitly measure the time taken to setup the matrix equation before the

first time step. This constituted a small fraction of the total run time even for a few time step, and remains constant as the number of time steps increase. There is some overlap in the time spent assembling the matrix equation and processing calls to the framework, because the assembly method calls the framework at various points. In the tests, we used 1500 finite element cells along the cochlear partition, resulting in 3000 radial segments. This is about one segment per inner hair cell – essentially the largest number of radial segments that might be used in a cochlear model. Typical results showed that solving the matrix equation dominated the simulation, justifying our focus on the solver in the remainder of the tests.

We tested two ways of using UMFPack: reinitialising the solver (recomputing the LU factorisation) at each time step; and initialising the solver just once, before the first time step. Three variants of the CG solver were also tested: the deal.II CG solver with an SSOR preconditioner; the deal.II CG without a preconditioner; and the Trilinos CG solver with an SSOR preconditioner. The SSOR relaxation parameter was set to 1.8 in all cases. In order to only measure the solver performance, we deliberately excluded the data collection process required in the case of the Trilinos solver from our measurements of solver time. This means that our measurements for the MPI solver show the best performance that could be expected from the solver. The reported results for the CG solvers used an error bound of 1×10^{-15} . (See the deal.II documentation, [Bangerth and Kanschat 2011](#), for an explanation of this bound.)

Tests were performed on two platforms: an 8 core node consisting of two quad-core 1.60 GHz E5310 Intel[®] Xeon[®] processors; and two 4 core nodes consisting of two dual-core 2.20 GHz model 275 AMD Opteron[™] processors.

B.3 Results

When UMFPack was used as the solver, and the solver object was initialised each time step (requiring an LU factorisation of \mathbf{K}_{aug}), solving the system took 92% of the run-time. Assembling the matrix equation took 5.3% of the run-time, and calls to the framework required just 2.8% of the run-time. Therefore, our priority in any optimisation must be to reduce the time taken to solve the system.

Figure B.1 on the following page shows results from the tests that gave a speedup with increasing numbers of cores, with some of the tests that were essentially serial included for comparison. The two tests that used the Trilinos CG solver with an SSOR preconditioner showed good speed up over the full range tested (8 processors). In the case of the Opteron nodes, MPI threads 5 to 8 run on cores that are on a separate motherboard so communication between these nodes goes over the Ethernet network, but the nearly linear speedup is still observed. The two tests using just the deal.II CG solver without a preconditioner showed a slight speedup (far from linear however). While the




Figure B.1: A comparison of time to solve the linear system using various processors and solvers. SSOR MPI uses the Trilinos implementation of symmetrical successive over relaxation and CG. SSOR deal.II uses the deal.II implementation of SSOR and CG. The remaining trends marked “deal.II” use the deal.II CG solver without preconditioning. UMFPack marks the performance of the UMFPack solver when it is reinitialised at each time step. The trend marked “linear speedup” shows the slope that results from the run-time halving each time the number of cores doubles. Trends marked “Xeon” used 1.60 GHz Xeon processors. Trends marked “Opteron” used 2.20 GHz Opteron processors.

speedup observed in these tests is fairly poor, it has been obtained without any extra effort to convert the serial code to parallel code.

When the UMFPack solver and the deal.II CG solver with SSOR preconditioner were used, we did not observe any significant speedup. However, these methods are about twice as fast when one core is used for the Trilinos solver, and remain competitive until 3 MPI threads are used. The significance of this is that, given sufficient RAM on the nodes, the serial solvers can perform three tests in parallel, and still obtain similar performance to the multithreaded solver for each test. Therefore, parallel solvers are outperformed by serial solvers, for this particular problem, and the using the solvers and preconditioners considered.

Results from the Xeon processor shows the CG solver with SSOR in deal.II outperforming UMFPack slightly (by 9.6 s per time step versus 12.3 s per time step). If the UMFPack solver is only initialised once, the UMFPack solver takes just 0.7 s per time step. This order of magnitude improvement can only be obtained if \mathbf{K}_{aug} remains the same for all time steps.

B.4 Discussion

The results showed that LU factorisation and CG with a full SSOR preconditioner, which are both essentially serial algorithms, outperform the parallel algorithms. This is largely due to the size of the problem (which has approximately 90000 unknowns). LU factorisation and the deal.II SSOR preconditioner both operate on the whole matrix, and the operations would become increasingly prohibitive for larger problems. However, we are already considering a large 2D cochlear model, so we would need to go to a 3D model, or a different fluid model before the matrix size would increase significantly.

The fact that the UMFPack solver with reinitialisation each time step is outperformed by the deal.II CG solver with an SSOR preconditioner is largely due to the error bound selected. The tolerance of 1×10^{-15} was selected to ensure there is no visible difference between the results from the two methods when plotted, for the small number of time steps that were tested. The competitive results suggest that both of these solvers should be considered as options for full length simulations.

One of the surprises in the data presented in figure B.1 is that the Xeon processors perform as well as, or even outperform, the Opteron processors while running the same code. The Opteron processors have both a higher CPU speed and faster access to RAM, so they should run the jobs faster than the Xeon processors. A possible explanation could be that the size of the processor cache is affecting the results. The Xeon processors have 4096 kB of cache to the Opteron's 1024 kB cache.

In this appendix, we have studied the performance of large 2D cochlear models implemented using this framework. The solver time dominates the overall run-time, with calls to the framework and assembly of the matrix equation taking approximately 2.8% and 5.3% of the run-time. Therefore, any efforts to reduce the run-time of the code must be directed at the solver. Careful parallelisation of the entire code using the Trilinos library could yield nearly linear speedup. However, the two serial solvers investigated were twice as fast as the parallel solvers on a single core. The parallel solvers began to break-even for around three cores. This shows that it is more efficient to run multiple jobs in parallel using a serial solver, rather than one multithreaded job after another. (For these problems, at this point in time. Future hardware and algorithms may cause the picture to change.) Sparse direct methods, such as UMFPack, are known to work well on 2D problems with fewer than 100000 unknowns, and it was the solver of choice for most of the results presented in this thesis.

Appendix C

Unification and Extension of Monolithic State Space and Iterative Cochlear Models

This appendix reproduces the article [Rapson *et al.* \(2012\)](#), published in the Journal of the Acoustical Society of America. The text presented represents the ASA-formatted version, as the journal kindly permits. My coauthors, David Karpul and Jonathan Tapson, gave their consent for me to include the article verbatim. Although their assistance with the paper is much appreciated, it represents my own work for the purposes of assessment of this thesis.

Unification and extension of monolithic state space and iterative cochlear models

Michael J. Rapson^{a)} and Jonathan C. Tapson

Department of Electrical Engineering, University of Cape Town, Private Bag, Rondebosch, 7701, Cape Town, South Africa

David Karpul

UCT/MRC Research Unit for Exercise Science and Sports Medicine, Department of Human Biology, Faculty of Health Sciences, University of Cape Town, 3rd Floor, Sports Science Institute, Boundary Road, Newlands, 7725, Cape Town, South Africa

(Received 14 April 2011; revised 2 December 2011; accepted 9 March 2012)

Time domain cochlear models have primarily followed a method introduced by Allen and Sondhi [J. Acoust. Soc. Am. **66**, 123–132 (1979)]. Recently the “state space formalism” proposed by Elliott *et al.* [J. Acoust. Soc. Am. **122**, 2759–2771 (2007)] has been used to simulate a wide range of non-linear cochlear models. It used a one-dimensional approach that is extended to two dimensions in this paper, using the finite element method. The recently developed “state space formalism” in fact shares a close relationship to the earlier approach. Working from Diependaal *et al.* [J. Acoust. Soc. Am. **82**, 1655–1666 (1987)] the two approaches are compared and the relationship formalized. Understanding this relationship allows models to be converted from one to the other in order to utilize each of their strengths. A second method to derive the state space matrices required for the “state space formalism” is also presented. This method offers improved numerical properties because it uses the information available about the model more effectively. Numerical results support the claims regarding fluid dimension and the underlying similarity of the two approaches. Finally, the recent advances in the state space formalism [Bertaccini and Sisto, J. Comp. Phys. **230**, 2575–2587 (2011)] are discussed in terms of this relationship.

© 2012 Acoustical Society of America. [http://dx.doi.org/10.1121/1.3699238]

PACS number(s): 43.64.Kc, 43.64.Bt [KG]

Pages: 3935–3952

I. INTRODUCTION

Computational cochlear models are widely used to test theories relating to cochlear physiology. Time-domain cochlear models have been applied to give insight in a range of hearing phenomena including otoacoustic emissions (van Hengel, 1996; van Hengel *et al.*, 1996; Moleti *et al.*, 2009) and distortion products (Matthews, 1980; Kim *et al.*, 1980; Liu and Neely, 2010; How *et al.*, 2010). They have proved their worth as research tools in a field where physiological data is difficult to obtain due to fragility and inaccessibility of the cochlea.

Two major algorithms appear in the literature for solving computational cochlear models in the time domain. These approaches appear very different as presented in the literature, but we show here that they share a close relationship. Understanding this relationship allows researchers to leverage the advantages that each method offers and to convert models between the two forms. The second (more recent) algorithm has only been implemented in one-dimensional (1D) models. We also show how the finite element method (FEM) may be used to extend this algorithm to two-dimensional (2D) models.

Allen and Sondhi (1979) introduced the first time domain cochlear models and their approach was developed further by Diependaal *et al.* (1987). A number of groups have

used similar methods to make 1D and 2D time domain models. (Allen and Sondhi, 1979; Matthews, 1980; Diependaal and Viergever, 1989, are examples of 2D implementations.) We describe this approach as *monolithic iterative* (MI). The word iterative refers to the fact that the computation to advance the simulation by one time step involves a series of stages.¹ Global system matrices are not formed in this method, but in the course of clarifying the relationship between the two algorithms it will become clear that equivalent global matrices exist, hence the method is implicitly monolithic.

More recently, Elliott *et al.* (2007) introduced an alternative approach based on certain state space techniques. Their technique has been referred to as a “state space model.” We prefer to designate it as *monolithic state space* (MSS).² In this case the global matrices for the monolithic method are explicitly formed.

The MSS method, in turn, has been developed further by Bertaccini and Sisto (2011). A major benefit of this approach is that it allows the eigenvalues (which show stability) of a linear cochlear model to be computed easily. These methods use finite difference approximations for spatial discretization, and have all been implemented as 1D models. To date finite differences (FD) have been used to spatially discretize the cochlear fluid in these models.

As we have noted these two methods in fact share a close relationship. By expressing the MI method in the notation of

^{a)} Author to whom correspondence should be addressed. Electronic mail: rpsmic001@gmail.com

the MSS method we derive an alternative approach to assembling the state space matrices required. We refer to this approach as MSS *variant* or MSSV. It can also be used to compute the model's eigenvalues or compute the time domain response, and has certain desirable numerical properties. Ultimately, however, our aim is not to argue that any of the three approaches discussed is superior, but to note their specific strengths and weaknesses. We derive an identity that not only allows the differences between the methods to be summarized, but also offers insight into the numerical stiffness problems mentioned in Sisto *et al.* (2010) and Bertaccini and Sisto (2011). Since the MSS method is a relative newcomer to cochlear modeling, we believe that describing its relationship to the established approach will assist users of cochlear models to understand their options.

We present our findings using the finite element method (FEM) and a 2D fluid model. The FEM matrices derived here allow the theory developed in Elliott *et al.* (2007) to be applied without change to models with a 2D fluid. They also allow greater freedom in the fluid geometry. The literature suggests that these changes will improve the accuracy of the cochlear model. Janssen *et al.* (1978) were the first to apply the finite element method to cochlear modeling [see also Viergever (1980)]. They noted how it readily allows local refinement of the spatial mesh as one of its benefits. Kagawa *et al.* (1987) solved both 2D and 3D cochlear models using the finite element method and used the model to investigate the spiral shape of the cochlea.

A number of authors have noted that a 2D fluid offers improved accuracy over 1D models (Steele and Taber, 1979; Viergever, 1980; Diependaal and Viergever, 1989; de Boer, 1991, section 2.3; Allen, 2001, pp. 14–15). There is a further improvement for 3D models but it is smaller than for the 1D to 2D improvement. Under some circumstances the accuracy of a 2D model is maintained if the second dimension is eliminated, for example by applying the Green's function method (Allen and Sondhi, 1979; Diependaal and Viergever, 1989). This is the case in the method we propose, as the numerical results show.

In the remainder of the paper we will first discuss the problem setting for the cochlear models to which these results apply. Secondly we will present a mathematical derivation of the results we mention above, followed by numerical results to support our claims. Finally we will discuss some of the implications of our findings. Nonlinear models and the recent work of Bertaccini and Sisto (2011) are considered at this stage.

The mathematical derivation is split into three parts. Firstly, we derive the FEM matrices required to extend the 1D FD model to 2D in Sec. III. In Sec. IV we express the MI method in terms of the MSS method to obtain the MSSV method and finally manipulate the expression for the MSSV method to show that it is related to the MSS method by an identity. An analysis of the numerical results presented is provided in Sec. V.

II. RELEVANT CLASSES OF COCHLEAR MODELS

The two broad modeling approaches discussed in this paper make some common assumptions about the cochlea.

Two fundamental assumptions are made: that the cochlear fluid is incompressible, irrotational and inviscid, allowing the behavior of the fluid to be described by the Laplace equation, $\nabla^2 P = 0$, where P is a scalar pressure field and ∇^2 is the Laplacian operator; and that the cochlear partition can be divided into a series of radial segments [called *elements* in Elliott *et al.* (2007)] that can each be described by a transfer function or state space model. A number of other assumptions are useful for simplifying the model's description and can more easily be removed. For example, it is common to assume that: the geometry of the cochlea is rectangular; the scala vestibuli and scala tympani are symmetrical allowing a single scala to be modeled; and a 1D or 2D representation of the cochlear fluid is adequate instead of its actual 3D structure. Boundary conditions are also required for the stapes and helicotrema (or round window, if the scala vestibuli and scala tympani are modeled separately). The results of this paper are relevant to any model that makes the first two assumptions. We adopt the other assumptions discussed in order to allow a clear derivation of the result.

III. EXTENDING THE MSS METHOD TO 2D

We first show that the MSS method in Elliott *et al.* (2007) can be extended to a 2D (or 3D) model by applying FEM. Most significantly, this extension does not change the form of the equations that they suggest, so the remainder of the theory developed continues to apply to the new model.

In order to show this, we first recall their notation and problem setting. \mathbf{A}_E , \mathbf{B}_E and \mathbf{C}_E are *combined state space* matrices. They have a block diagonal structure where each block is the state space matrix for an element of the cochlear partition. The first and last block incorporate the boundary conditions, the stapes and the helicotrema, respectively. They define a combined states vector $\mathbf{x}(t)$ and an acceleration vector $\ddot{\mathbf{w}}(t)$ with a similar ordering. Again these vectors consist mainly of cochlear partition variables, but also include boundary conditions. The pressures $\mathbf{p}(t)$ in a spatially discretized fluid are described by $\mathbf{p}(t) = \mathbf{F}^{-1}\ddot{\mathbf{w}}(t) + \mathbf{q}(t)$, where \mathbf{F} is a FD fluid matrix and $\ddot{\mathbf{w}}(t) + \mathbf{q}(t)$ is the total acceleration of the fluid boundary. They separate the stapes acceleration into a component due to the internal pressure, which is included in $\ddot{\mathbf{w}}(t)$, and an external forcing acceleration at the stapes, $\mathbf{q}(t)$.

Using these definitions, they show that a combined state space system can be formed:

$$\dot{\mathbf{x}}(t) = \mathbf{A}\mathbf{x}(t) + \mathbf{B}\mathbf{u}(t), \quad (1)$$

$$\ddot{\mathbf{w}}(t) = \mathbf{C}_E\dot{\mathbf{x}}(t), \quad (2)$$

where

$$\mathbf{A} = [\mathbf{I} - \mathbf{B}_E\mathbf{F}^{-1}\mathbf{C}_E]^{-1}\mathbf{A}_E, \quad (3)$$

$$\mathbf{B} = [\mathbf{I} - \mathbf{B}_E\mathbf{F}^{-1}\mathbf{C}_E]^{-1}\mathbf{B}_E, \quad (4)$$

$$\mathbf{u}(t) = \mathbf{F}^{-1}\mathbf{q}(t), \quad (5)$$

$\dot{\mathbf{x}}(t)$ is the time derivative of the states, $\mathbf{u}(t)$ is the input to the state update Eq. (1) and \mathbf{I} is the identity matrix. At this stage it

is clear that the single matrix Eq. (1) represents the dynamics of the whole system, including the cochlear partition, fluid and boundary conditions. Therefore the term “monolithic” as used in the fluid-structure interaction field seems to be an appropriate description of the approach.

In order to replace \mathbf{F}^{-1} in the above equations we work through the finite element method, starting by describing partial differential equations that must be satisfied. The 1D partial differential equations used in Elliott *et al.* (2007) can be converted to their 2D equivalents. We require four equations. The first is $\nabla^2 p(x, y, t) = 0$, the Laplace equation. The second and third are momentum equations, stating that the pressure gradient normal to the fluid boundary at the point (x, y) , $\nabla p(x, y, t) \cdot \vec{n}$, is proportional to the normal acceleration at that point. Finally zero pressure at the helicotrema is enforced. The four equations are

$$\nabla^2 p(x, y, t) = 0 \text{ on } \Omega, \quad (6)$$

$$\nabla p(x, y, t) \cdot \vec{n} = -2\rho\ddot{w}_{cp}(x, y, t) \text{ on } \Gamma_{cp}, \quad (7)$$

$$\nabla p(x, y, t) \cdot \vec{n} = -2\rho\ddot{w}_s(x, y, t) \text{ on } \Gamma_s, \quad (8)$$

$$p(x, y, t) = 0 \text{ on } \Gamma_H. \quad (9)$$

In these equations $p(x, y, t)$ is the pressure field, ρ is the fluid density, ∇ is the gradient operator and ∇^2 the Laplacian. Ω is the fluid region and \vec{n} is an outward unit normal. Γ is the boundary of the fluid region and it is divided into the cochlear partition Γ_{cp} , stapes Γ_s , helicotrema Γ_H and bony cochlear wall where no normal acceleration is allowed. Following the notation in Elliott *et al.* (2007), $\ddot{w}_{cp}(x, y, t)$ and $\ddot{w}_s(x, y, t)$ are the acceleration of the cochlear partition and stapes, respectively. The use of subscripts of cp and s in this manner will be followed generally.

The weak form of these equations can be found in the usual manner [see Fish and Belytschko (2007) for example]. We require a functional analysis space for a trial solution, $U = \{p(x, y, t) | p(x, y, t) \in H^1, p(x, y, t) = 0 \text{ on } \Gamma_H\}$ and a weighting function, $U_0 = \{q(x, y, t) | q(x, y, t) \in H^1, q(x, y, t) = 0 \text{ on } \Gamma_H\}$. H^1 is the functional analysis space where both the function and its derivative are square integrable. Using these spaces the weak form is

Find $p(x, y, t) \in U$ such that

$$\begin{aligned} & \int_{\Omega} (\nabla q(x, y, t))^T \nabla p(x, y, t) d\Omega \\ &= - \int_{\Gamma_{cp}} q(x, y, t)^T 2\rho\ddot{w}_{cp}(t) d\Gamma_{cp} \\ & - \int_{\Gamma_s} q(x, y, t)^T 2\rho\ddot{w}_s(t) d\Gamma_s \quad \forall w \in U_0, \end{aligned} \quad (10)$$

where \forall means “for all” and T is the transpose operator, which will allow us to write inner products as matrix multiplications.

We discretize the model according to FEM, by letting the continuous spaces be replaced by discrete spaces such that

$$p(x, y, t) \approx p^h(x, y, t) = \mathbf{N}(x, y)\mathbf{p}_n(t), \quad (11)$$

$$q(x, y, t) \approx q^h(x, y, t) = \mathbf{N}(x, y)\mathbf{q}_n(t), \quad (12)$$

$$\nabla p(x, y, t) \approx \nabla p^h(x, y, t) = \mathbf{B}(x, y)\mathbf{p}_n(t), \quad (13)$$

$$\nabla q(x, y, t) \approx \nabla q^h(x, y, t) = \mathbf{B}(x, y)\mathbf{q}_n(t), \quad (14)$$

where $\mathbf{N}(x, y) \in H^1$ are shape functions and $\mathbf{B}(x, y)$ are gradients of shape functions. These are both $1 \times n$ vectors, where n is the number of FEM nodes in the domain. $\mathbf{p}_n(t)$ and $\mathbf{q}_n(t)$ are the nodal pressures and weighting functions at a time t , respectively, $n \times 1$.

Finally, let $\mathbf{q}_F(t)$ be the subset of values in $\mathbf{q}_n(t)$ not on Γ_H [on Γ_H the value of $\mathbf{p}_n(t)$ is enforced separately]. Then the weak form becomes

$$\begin{aligned} & \mathbf{q}_n^T(t) \int_{\Omega} \mathbf{B}^T(x, y) \mathbf{B}(x, y) d\Omega \mathbf{p}_n(t) \\ & + \mathbf{q}_n^T(t) \int_{\Gamma_{cp}} \mathbf{N}^T(x, y) 2\rho\ddot{w}_{cp}(t) d\Gamma_{cp} \\ & + \mathbf{q}_n^T(t) \int_{\Gamma_s} \mathbf{N}^T(x, y) 2\rho\ddot{w}_s(t) d\Gamma_s = 0 \quad \forall \mathbf{q}_F(t). \end{aligned} \quad (15)$$

Gauss quadrature can be used to replace the integration by a summation over terms located at quadrature points. Gauss quadrature over Ω proceeds in the usual fashion, using a method appropriate for the element chosen, but integration over Γ is manipulated to allow the normal accelerations to be unknowns. Since Γ is a 1D line in a 2D space we introduce mapping from $\mathbb{R} \rightarrow \mathbb{R}^2$, $(x, y) = f(\xi)$ to describe locations on the boundary.

Consider the boundary integral $\mathbf{I}(t)$ over the cochlear partition

$$\mathbf{I}(t) = \int_{\Gamma_{cp}} \mathbf{N}^T(f(\xi)) 2\rho\ddot{w}_{cp}(t) d\Gamma_{cp}. \quad (16)$$

The Gauss quadrature rule allows the integration to be computed with

$$\mathbf{I}(t) = \sum_{i=1}^N W_i |\mathbf{J}(f(\xi_i))| \mathbf{N}^T(f(\xi_i)) 2\rho\ddot{w}_{cp_i}(t), \quad (17)$$

where N is the number of quadrature points on the cochlear partition in this section (n_{cp} will be used generally), W_i is the weighting of the i th quadrature point and $\mathbf{N}^T(f(\xi_i)) 2\rho\ddot{w}_{cp_i}$ is a vector of contributions to the nodal values $\mathbf{I}(t)$ due to the function at the i th quadrature point. $|\mathbf{J}(f(\xi_i))|$ accounts for the mapping between Γ_{cp} and the parent domain where the integration is performed. For compactness we do not show the mapping f explicitly in future equations, so expressions like $\mathbf{N}(f(\xi))$ are written $\mathbf{N}(\xi)$. Then Eq. (17) can be written as

$$\mathbf{I} = 2\rho\mathcal{W}_{cp}\ddot{w}_{cp}(t), \quad (18)$$

where

$$\begin{aligned} \mathcal{W}_{cp} = & [W_1 |\mathbf{J}(\xi_1)| \mathbf{N}^T(\xi_1) W_2 |\mathbf{J}(\xi_2)| \mathbf{N}^T(\xi_2) \\ & \cdots W_N |\mathbf{J}(\xi_N)| \mathbf{N}^T(\xi_N)], \end{aligned}$$

$$\ddot{\mathbf{w}}_{cp}(t) = [\ddot{w}_{cp_1}(t)\ddot{w}_{cp_2}(t)\dots\ddot{w}_{cp_n}(t)]^T.$$

This procedure can be repeated with the stapes integral over Γ_s .

Applying the Gauss quadrature rule above to the boundary terms and assuming that a method is used to enforce the boundary condition Γ_H , Eq. (15) can be written as

$$\left\{ \int_{\Omega} \mathbf{B}(x,y)\mathbf{B}(x,y)d\Omega \right\} \mathbf{p}_n(t) + \{2\rho\mathcal{W}_{cp}\} \ddot{\mathbf{w}}_{cp}(t) + \{2\rho\mathcal{W}_s\} \ddot{\mathbf{w}}_s(t) = 0.$$

This matrix equation will be written as

$$\mathbf{K}\mathbf{p}_n(t) - \mathbf{R}_{cp}\ddot{\mathbf{w}}_{cp}(t) - \mathbf{R}_s\ddot{\mathbf{w}}_s(t) = 0, \quad (19)$$

where

$$\mathbf{K} = \int_{\Omega} \mathbf{B}(x,y)\mathbf{B}(x,y)d\Omega, \quad (20)$$

$$\mathbf{R}_{cp} = -2\rho\mathcal{W}_{cp}, \quad (21)$$

$$\mathbf{R}_s = -2\rho\mathcal{W}_s. \quad (22)$$

These matrices are of size $n \times n$, $n \times n_{cp}$ and $n \times n_s$, respectively, and terms in \mathcal{W}_{cp} and \mathcal{W}_s are of the form $W_i|J(\xi_i)|N(\xi_i)$, which is $n \times 1$ for fixed i .

The pressures $\mathbf{p}_n(t)$ are located at the nodes. These can be interpolated to the quadrature points on the cochlear partition and the stapes by using the shape functions, $\mathbf{N}(x,y)$, and the quadrature point locations. Let these pressures be $\mathbf{p}_{cp}(t)$ and $\mathbf{p}_s(t)$, respectively. Then,

$$\mathbf{p}_{cp}(t) = \mathbf{Q}_{cp}\mathbf{p}_n(t), \quad (23)$$

$$\mathbf{p}_s(t) = \mathbf{Q}_s\mathbf{p}_n(t), \quad (24)$$

where

$$\mathbf{Q}_{cp} = \left[\mathbf{N}^T(\xi_1)\mathbf{N}^T(\xi_2)\dots\mathbf{N}^T(\xi_{n_{cp}}) \right]^T \mathbf{p}_n(t),$$

$$\mathbf{Q}_s = \left[\mathbf{N}^T(\varepsilon_1)\mathbf{N}^T(\varepsilon_2)\dots\mathbf{N}^T(\varepsilon_{n_s}) \right]^T \mathbf{p}_n(t).$$

ξ_i and ε_i are quadrature point locations on the cochlear partition and stapes, respectively. \mathbf{Q}_{cp} and \mathbf{Q}_s are $n_{cp} \times n$ and $n_s \times n$, respectively.

We are now ready to form \mathbf{F}^{-1} using the FEM variables that we have defined. The resulting matrix must have the same order of $\mathbf{x}(t)$ as previously. We therefore define

$$\mathbf{R} = [\mathbf{R}_s \quad \mathbf{R}_{cp}], \quad (25)$$

$$\mathbf{Q}^T = \left[\mathbf{Q}_s^T \quad \mathbf{Q}_{cp}^T \right], \quad (26)$$

$$\ddot{\mathbf{w}}^T(t) = \left[\ddot{\mathbf{w}}_s^T(t) \quad \ddot{\mathbf{w}}_{cp}^T(t) \right], \quad (27)$$

$$\mathbf{p}(t) = \left[\mathbf{p}_s(t) \quad \mathbf{p}_{cp}(t) \right]. \quad (28)$$

Note that \mathbf{R}_s and \mathbf{Q}_s include terms for the n_s quadrature points on the stapes, whereas a 1D model offers only one

value on the stapes. There are various ways to use this additional flexibility, but for comparison with the 1D model we elect to constrain the stapes acceleration to be spatially constant. This assumes that the stapes moves like a piston, as in the 1D model. To enforce this we substitute $\ddot{\mathbf{w}}_s = [1 \ 1 \ \dots \ 1]^T \ddot{w}_s$, where \ddot{w}_s is a single value. Therefore, \mathbf{R}_s can be replaced by $\bar{\mathbf{R}}_s = \mathbf{R}_s [1 \ 1 \ \dots \ 1]^T$ in the expression for \mathbf{R} . Elliott *et al.* (2007) separate the stapes acceleration into a forcing acceleration and a component dependent on the pressure at the stapes. If this is done again then $w_1 + q_1 = \ddot{w}_s$. A similar process can be applied to \mathbf{Q}_s . In this case we assume that the integral of the pressure over the stapes causes the force of the fluid acting on it. Taking this into account, we find $\bar{\mathbf{Q}}_s$ that is $1 \times n$ to replace the original \mathbf{Q}_s in (26). See Appendix A for a derivation of $\bar{\mathbf{Q}}_s$.

Using Eqs. (19) and (23)–(28) we obtain

$$\mathbf{p}(t) = \mathbf{Q}\mathbf{K}^{-1}\mathbf{R}\ddot{\mathbf{w}}(t) + \mathbf{Q}\mathbf{K}^{-1}\mathbf{R}\mathbf{q}(t), \quad (29)$$

where we have used $\mathbf{q}(t)$ to represent the stapes driving acceleration as in Elliott *et al.* (2007).

By comparing this expression and Eq. (24) in Elliott *et al.* (2007) we find that

$$\mathbf{F}^{-1} = \mathbf{Q}\mathbf{K}^{-1}\mathbf{R}. \quad (30)$$

This definition for \mathbf{F}^{-1} can be used in Eqs. (3)–(5) above, extending the MSS method to 2D domains and FEM. The shape of the domain need not be a rectangle, and the extension to 3D follows along similar lines. FEM packages, such as deal.II (Bangerth *et al.*, 2007, 2008), used in this work, simplify the computation of the required finite element matrices \mathbf{Q} , \mathbf{K} , and \mathbf{R} , by providing the necessary shape functions, Gauss quadrature rules and quadrature point locations.

IV. EQUIVALENCE OF STATE SPACE METHOD TO PREVIOUS APPROACHES

Allen and Sondhi introduced the first time domain cochlear model in 1979 (Allen and Sondhi, 1979). They used a Green's function to spatially discretize the cochlear fluid, and ensured stability of the time integration by defining an augmented Green's function kernel for the fluid. This kernel included contributions from the cochlear partition mass. Diependaal *et al.* investigated the time domain solution further using the method of Galerkin (Diependaal *et al.*, 1987), which is closely related to the finite element method. For convenience, the work of Diependaal *et al.* is selected as the starting point for a comparison between Elliott *et al.* (2007) and previous approaches. This requires two stages: firstly expressing their algorithm in the notation above, and secondly manipulating the resulting expression to show equivalence.

A. Expressing the algorithm in this notation

Once again we begin by recalling some expressions from the relevant paper, in this case Diependaal *et al.* (1987):

$$p''(x,t) - [2\rho\beta(x)/a(x)]\ddot{u}(x,t) = 0, \quad 0 < x < l, \quad t > 0, \quad (31)$$

$$p'(0,t) = f(t), \quad t \geq 0, \quad (32)$$

$$p(x,t) = m(x)\ddot{u}(x,t) + r(x,t)\dot{u}(x,t) + s(x,t)u(x,t). \quad (33)$$

The first two equations refer to the fluid and stapes boundary condition. It is a 1D formulation that is identical to the one used in Elliott *et al.* (2007) except for the factor relating to the basilar membrane width $\beta(x)$ that Diependaal includes. The term $a(x)$ is the cochlear height, designated H in Elliott *et al.* (2007). Diependaal *et al.* (1987) studied a specific form for the elements [as opposed to the general state space structure considered in Elliott *et al.* (2007)]. In Diependaal *et al.*, $m(x)$, $r(x,t)$ and $s(x,t)$ are the mass, damping and stiffness of the element, respectively. $u(x,t)$ is the displacement of a particular element. Considering the stapes boundary condition we find that it is the same except for a weighting, that is $2\rho\dot{w}_s(t) = f(t)$. The final equation is the transfer function for cochlear partition elements. Diependaal *et al.* define two further variables related to the cochlear partition:

$$\alpha^2(x) = 2\rho\beta(x)/m(x)a(x), \quad (34)$$

$$g(x,t) = r(x,t)\dot{u}(x,t) + s(x,t)u(x,t). \quad (35)$$

Then

$$p''(x,t) - \alpha^2(x)p(x,t) = -\alpha^2(x)g(x,t). \quad (36)$$

These are Eqs. (1), (6), (2), and (7)–(9) in Diependaal *et al.* (1987).

We now compare the equations in both papers to observe equivalent expressions. Starting with the cochlear partition we find

$$\ddot{w}_n(t) = \mathbf{C}_n\mathbf{A}_n\mathbf{x}_n(t) + \mathbf{C}_n\mathbf{B}_n p_n(t), \quad (37)$$

$$\ddot{u}(x,t) = \frac{1}{m(x)}p(x,t) - \left(\frac{r(x,t)}{m(x)}\dot{u}(x,t) + \frac{s(x,t)}{m(x)}u(x,t) \right), \quad (38)$$

where the notation with n as a subscript refers to individual cochlear partition elements. $\ddot{w}_n(t) = \ddot{u}(x,t)$ are the accelerations and $\mathbf{x}_n(t) = [\dot{u}(x,t)u(x,t)]$ are the states. Therefore,

$$\mathbf{C}_n\mathbf{B}_n = \frac{1}{m(x)}, \quad (39)$$

$$\mathbf{C}_n\mathbf{A}_n = \left[-\frac{r(x,t)}{m(x)} - \frac{s(x,t)}{m(x)} \right]. \quad (40)$$

From Eqs. (34) and (35) we see that

$$\alpha^2(x_n) = 2\rho\mathbf{C}_n\mathbf{B}_n, \quad (41)$$

$$g(x_n,t) = -(\mathbf{C}_n\mathbf{B}_n)^{-1}\mathbf{C}_n\mathbf{A}_n\mathbf{x}_n(t). \quad (42)$$

The term H which would normally appear in the expression for $\alpha^2(x_n)$ is not required in a 2D formulation because the

integration over the cochlear height, performed at a later stage, accounts for it.

Diependaal *et al.* form a matrix equation $\mathbf{A}_d\mathbf{p}(t) = \mathbf{k}(t)$, by collecting terms of $g(x,t)$ and $\alpha^2(x)$ as well as considering the differential equation and its boundary conditions. From these definitions and Eq. (15) in Diependaal *et al.* (1987) the reader can verify that

$$\mathbf{k}(t) = \mathbf{RC}_E\mathbf{A}_E\mathbf{x}(t) + \mathbf{Rq}(t), \quad (43)$$

where \mathbf{R} accounts for the factor of 2ρ and the integration. The effect of $f(t)$ is shared between $\mathbf{q}(t)$ and terms in $\mathbf{x}(t)$ on the stapes. The change from subscripts n to E indicates that matrices for the elements of the cochlear partition have been replaced by their combined equivalents.

The matrix \mathbf{A}_d contains the terms $p''(x,t) - \alpha^2(x)p(x,t)$. Consider Eq. (19) above. If the definitions in (25)–(28) are used, the combined form of (37) is substituted, and we convert from pressures at quadrature points to pressures at nodes using $\mathbf{p}(t) = \mathbf{Qp}_n(t)$, we obtain

$$\mathbf{Kp}_n(t) - \mathbf{RC}_E\mathbf{B}_E\mathbf{Qp}_n(t) = \mathbf{RC}_E\mathbf{A}_E\mathbf{x}(t) + \mathbf{Rq}(t). \quad (44)$$

Since the right hand side is $\mathbf{k}(t)$ the left hand side shows

$$\mathbf{A}_d = \mathbf{K} - \mathbf{RC}_E\mathbf{B}_E\mathbf{Q}. \quad (45)$$

Finally Diependaal *et al.* express the element acceleration, called $\ddot{\mathbf{w}}(t)$ in this paper, as

$$\dot{\mathbf{v}}(t) = \mathbf{Q}[\mathbf{p}(t) - \mathbf{g}(t)]. \quad (46)$$

If we extend this equation to include the stapes, then $\mathbf{Q} = \mathbf{C}_E\mathbf{B}_E$ and this is equivalent to

$$\ddot{\mathbf{w}}(t) = \mathbf{C}_E\mathbf{A}_E\mathbf{x}(t) + \mathbf{C}_E\mathbf{B}_E\mathbf{Qp}_n(t), \quad (47)$$

the combined form of Eq. (37). Note that the total acceleration of the stapes would be found by adding the forcing acceleration: $\ddot{w}_s = \ddot{\mathbf{w}}_1 + \mathbf{q}_1$.

After preliminary steps, the algorithm that Diependaal *et al.* propose requires three steps to compute an estimate of $\ddot{\mathbf{w}}(t)$:

- (1) Compute $\mathbf{k}(t)$, from Eq. (43), which forms the right hand side of Eq. (44).
- (2) Solve for \mathbf{p} , solving our Eq. (44).
- (3) Compute $\mathbf{Q}(\mathbf{p} - \mathbf{g})$, which is Eq. (47).

Therefore, substituting Eq. (44) into (47) to eliminate $\mathbf{p}_n(t)$ yields

$$\ddot{\mathbf{w}}(t) = \mathbf{C}_E[\mathbf{I} + \mathbf{B}_E\mathbf{Q}[\mathbf{K} - \mathbf{RC}_E\mathbf{B}_E\mathbf{Q}]^{-1}\mathbf{RC}_E]\mathbf{A}_E\mathbf{x}(t) + \mathbf{C}_E\mathbf{B}_E\mathbf{Q}[\mathbf{K} - \mathbf{RC}_E\mathbf{B}_E\mathbf{Q}]^{-1}\mathbf{Rq}(t). \quad (48)$$

The complete (linked) state space system can be obtained by noting that $\ddot{\mathbf{w}}(t) = \mathbf{C}_E\dot{\mathbf{x}}(t)$ and that we have been working with the combined but unlinked state space equations through out. Therefore, it is correct to simply drop the initial \mathbf{C}_E matrix (which selects just the $\ddot{\mathbf{w}}(t)$ terms) to return to the complete system. Hence,

$$\mathbf{A} = [\mathbf{I} + \mathbf{B}_E \mathbf{Q} [\mathbf{K} - \mathbf{R} \mathbf{C}_E \mathbf{B}_E \mathbf{Q}]^{-1} \mathbf{R} \mathbf{C}_E] \mathbf{A}_E, \quad (49)$$

and

$$\mathbf{B} \mathbf{F}^{-1} = \mathbf{B}_E \mathbf{Q} [\mathbf{K} - \mathbf{R} \mathbf{C}_E \mathbf{B}_E \mathbf{Q}]^{-1} \mathbf{R}. \quad (50)$$

Based on the fact that these equations were derived from the same set of partial differential equations as those presented in Elliott *et al.* (2007), it is reasonable to expect that they are equivalent to those equations. We next show this equivalence by mathematical manipulation of the expressions and in the process we show the formal relationship between the expressions.

B. Manipulations to establish equality

Having expressed the MI method in the notation of Elliott *et al.* (2007) using a 2D FEM model, we manipulate Eqs. (49) and (50) to show that they are identical to Eqs. (1)–(5). In order to do this we require an identity for the inverse of a sum of matrices. We recall Eq. (2) in Henderson and Searle (1981),

$$(\mathbf{A}_h - \mathbf{U} \mathbf{D}^{-1} \mathbf{V})^{-1} = \mathbf{A}_h^{-1} + \mathbf{A}_h^{-1} \mathbf{U} (\mathbf{D} - \mathbf{V} \mathbf{A}_h^{-1} \mathbf{U})^{-1} \mathbf{V} \mathbf{A}_h^{-1}. \quad (51)$$

Matching the right hand side of (51) to (49) we find that $\mathbf{A}_h^{-1} = \mathbf{I}^{-1} = \mathbf{I}$, $\mathbf{U} = \mathbf{B}_E \mathbf{Q}$, $\mathbf{D} = \mathbf{K}$ and $\mathbf{V} = \mathbf{R} \mathbf{C}_E$. Therefore, after applying the identity,

$$\mathbf{A} = [\mathbf{I} - \mathbf{B}_E \mathbf{Q} \mathbf{K}^{-1} \mathbf{R} \mathbf{C}_E]^{-1} \mathbf{A}_E. \quad (52)$$

This matches Eq. (3), because $\mathbf{F}^{-1} = \mathbf{Q} \mathbf{K}^{-1} \mathbf{R}$.

We manipulate (50) so that we can select the same definitions of \mathbf{A}_h^{-1} , \mathbf{U} , \mathbf{D} , and \mathbf{V} . This requires us to introduce a new expression for $\mathbf{q}(t)$, $\mathbf{q}(t) = \mathbf{C}_E \tilde{\mathbf{q}}(t)$, and to add and subtract $\mathbf{I} \tilde{\mathbf{q}}(t)$ as follows:

$$\mathbf{B} \mathbf{F}^{-1} \mathbf{q}(t) = \mathbf{I}^{-1} \tilde{\mathbf{q}}(t) + \mathbf{B}_E \mathbf{Q} [\mathbf{K} - \mathbf{R} \mathbf{C}_E \mathbf{B}_E \mathbf{Q}]^{-1} \times \mathbf{R} \mathbf{C}_E \tilde{\mathbf{q}}(t) - \mathbf{I} \tilde{\mathbf{q}}(t). \quad (53)$$

Now applying the identity:

$$\mathbf{B} \mathbf{F}^{-1} \mathbf{q}(t) = [\mathbf{I} - \mathbf{B}_E \mathbf{Q} \mathbf{K}^{-1} \mathbf{R} \mathbf{C}_E]^{-1} \tilde{\mathbf{q}}(t) - \mathbf{I} \tilde{\mathbf{q}}(t). \quad (54)$$

Finally we note that

$$\mathbf{I} = [\mathbf{I} - \mathbf{B}_E \mathbf{Q} \mathbf{K}^{-1} \mathbf{R} \mathbf{C}_E]^{-1} [\mathbf{I} - \mathbf{B}_E \mathbf{Q} \mathbf{K}^{-1} \mathbf{R} \mathbf{C}_E]. \quad (55)$$

Expanding \mathbf{I} in this manner allows us to cancel a common $[\mathbf{I} - \mathbf{B}_E \mathbf{Q} \mathbf{K}^{-1} \mathbf{R} \mathbf{C}_E]^{-1} \tilde{\mathbf{q}}(t)$ term, to obtain the following expression:

$$\mathbf{B} \mathbf{F}^{-1} \mathbf{q}(t) = [\mathbf{I} - \mathbf{B}_E \mathbf{Q} \mathbf{K}^{-1} \mathbf{R} \mathbf{C}_E]^{-1} \mathbf{B}_E \mathbf{Q} \mathbf{K}^{-1} \mathbf{R} \mathbf{C}_E \tilde{\mathbf{q}}(t). \quad (56)$$

Since $\mathbf{F}^{-1} = \mathbf{Q} \mathbf{K}^{-1} \mathbf{R}$ and $\mathbf{q}(t) = \mathbf{C}_E \tilde{\mathbf{q}}(t)$, we have $\mathbf{B} = [\mathbf{I} - \mathbf{B}_E \mathbf{F}^{-1} \mathbf{C}_E]^{-1} \mathbf{B}_E$ as required.

Therefore when the MSS and MI methods are expressed in the same notation they appear very different, but the

resulting matrix expressions share a close relationship. Since they share this close relationship, it is clear that the algorithm proposed in Diependaal *et al.* (1987) is also a *monolithic* approach to modeling the cochlea. It is iterative in the sense that the original algorithm requires a series of steps to compute a new estimate of $\dot{\mathbf{x}}(t)$. We therefore refer to the original algorithm as *monolithic iterative* or MI. This section has shown that the algorithm can also be implemented by forming large state space matrices, similar to the MSS method. We will distinguish this alternative way of assembling the matrices by referring to the MSS *variation* (MSSV) method.

C. Summary

In this section we have shown that

$$\begin{aligned} \mathbf{A} &= [\mathbf{I} - \mathbf{B}_E \mathbf{Q} \mathbf{K}^{-1} \mathbf{R} \mathbf{C}_E]^{-1} \mathbf{A}_E \\ &= [\mathbf{I} + \mathbf{B}_E \mathbf{Q} [\mathbf{K} - \mathbf{R} \mathbf{C}_E \mathbf{B}_E \mathbf{Q}]^{-1} \mathbf{R} \mathbf{C}_E] \mathbf{A}_E, \end{aligned} \quad (57)$$

$$\begin{aligned} \mathbf{B} \mathbf{F}^{-1} &= [\mathbf{I} - \mathbf{B}_E \mathbf{Q} \mathbf{K}^{-1} \mathbf{R} \mathbf{C}_E]^{-1} \mathbf{B}_E \mathbf{Q} \mathbf{K}^{-1} \mathbf{R} \\ &= \mathbf{B}_E \mathbf{Q} [\mathbf{K} - \mathbf{R} \mathbf{C}_E \mathbf{B}_E \mathbf{Q}]^{-1} \mathbf{R}, \end{aligned} \quad (58)$$

where the first expression in each case is from the MSS method and the second is from the MSSV method. We will present computational results that suggest that these expressions are identical to the limit of numerical precision, and will discuss the strengths of each approach. In particular, the MSSV method has certain desirable numerical properties. The results will also show that the FEM implementation matches our expectations of a model with a 2D fluid.

V. RESULTS

In order to verify the claims made in this paper, four separate programs were written in C++ using the deal.II FEM library (Bangerth *et al.*, 2007, 2008). They all implemented the example used in Elliott *et al.* (2007), originally in Neely and Kim (1986), which uses data for the cat cochlea. These all share a common object class that provides element state space matrices given a location and value for γ , the active feedback gain, which guarantees that the parameters they receive are comparable. Table I provides a summary of the results which will be presented. It lists the claims we wish to support, and indicates which programs are used to support them.

The first code is a reimplemention of Elliott *et al.* (2007) (FD MSS) in C++. It is compared to the published results for the model and is then used as a benchmark to compare the other codes against.

The second, FE MSS is a FEM implementation of the model, using the matrix \mathbf{F}^{-1} discussed in Sec. III. FD MSS and FE MSS share a large amount of code. The method used to construct \mathbf{F}^{-1} is the major difference between the programs, but there are also small differences in element locations and these differences carry over into the methods that postprocess the results. Maintaining similarity between the programs ensures that any differences observed are due to the dimensional differences and are not artifacts of the

TABLE I. Summary of the code variations tested in this work, with a breakdown of their purposes. FD, finite difference; MSS, monolithic state space; FE, finite element; MSSV, monolithic state space variant; MI, monolithic iterative. FD MSS is an implementation of Elliott *et al.* (2007). FE MI follows the algorithm in Diependaal *et al.* (1987). FE MSS and FE MSSV are the two FEM extensions to Elliott *et al.* (2007) described in this paper. The text explains each test in more detail.

Purpose	FD MSS	FE MSS	FE MSSV	FE MI
Direct comparison with literature	used			
Test dimensional accuracy	used	used		
Discuss dimensional effects	used	used		
Numerical precision calculations	used	used	used	
Time domain simulations			all used	
Apical boundary condition tests			all used	

computation. Results from this pair of programs are therefore used to show that the proposed FEM matrices offer an extension to 2D accuracy. A 1D variant of this code was used to distinguish between the effects of changing from FD to FEM from the effects of using a 2D versus a 1D model.

FE MSSV is also a FEM implementation of the model. In this case the alternative method discussed in Sec. IV is used. There is a high degree of code similarity between FE MSS and FE MSSV. The two programs differ only in the way that the constituent matrices are combined to form the final state space matrices, **A** and **B**. Therefore, this pair of programs is used to support the claim that the MSSV method, developed from the Diependaal *et al.* MI method, is identical to the Elliott *et al.* MSS method, to numerical accuracy. Note that, as in Bertaccini and Sisto (2011), in the FE MSSV program the matrix \mathbf{F}^{-1} is not formed.

Finally, FE MI implements an algorithm along the lines of the iterative approach of Diependaal *et al.* (1987). In this case the final state space matrices **A** and **B** are not formed. The collected element matrices \mathbf{A}_E , \mathbf{B}_E and \mathbf{C}_E are also not required. This program only produces time domain results, and these are compared to time domain results from FE MSS and FE MSSV. Similarity in the time domain results show that both the MSS approach and its variant are equivalent to the MI approach of Diependaal *et al.* (1987).

Elliott *et al.* (2007) used 498 basilar membrane (BM) elements. The stapes and helicotrema boundaries each contribute an element towards the figure of 500 elements presented in that paper. Altogether there are 1995 states for that model. This same number of elements was used for the FD MSS program. The closest match to this for the remaining programs is to use 249 FEM cells along the BM, with two quadrature points per cell. This results in 498 BM elements and 1994 states including the stapes. There is no additional state from the helicotrema boundary condition because it is enforced directly on the fluid. In the height direction 10 FEM cells were used to ensure that the cells were square. The effect of using fewer cells in the vertical direction is discussed in Sec. V C.

A. Comparison to the literature

The program FD MSS produces results that are identical to previous publications for the model studied. Figure 5 from




FIG. 1. Plot of pressures vs distance from the stapes for five input frequencies. (a) The magnitude is expressed in dB re $P_e = 2 \times 10^{-5} \text{ Nm}^{-2}$ (rms) and (b) the phase is shown. Solid lines, FE MSS; dashed lines, FD MSS. Frequencies indicated on graph and further description is provided in the text. FE MSS shows characteristic 2D features. For example, compare this figure to Fig. 3 in Neely (1981).

Elliott *et al.* (2007) was reproduced from the final state space matrices obtained from FD MSS using SciPy Jones *et al.* (2001) and the expression for frequency response, Eq. (45) in Elliott *et al.* (2007). The stapes input acceleration was calculated from the middle ear transfer function provided in Neely and Kim (1986) assuming an input pressure at the eardrum of $2 \times 10^{-5} \text{ Nm}^{-2}$ (rms) at 1.6 kHz. In order to match the figure, it was found that $C_3(x) = 20e^{-80x} \text{ Nsm}^{-3}$, and $C_4(x) = 10400e^{-200x} \text{ Nsm}^{-3}$ from Neely and Kim (1986) were required. This provided a qualitative match, but to get the detail correct,³ the height H of the model needed to be 0.002 m. The other data required was as per Table I in Elliott *et al.* (2007). The eigenvalues of the **A** matrix (poles of the model) were plotted for $\gamma = 0$ and $\gamma = 1$. These plots agreed with Elliott *et al.* [2007, Fig. 7(a) and 7(b)] in terms of all visible features. These results confirm that we have correctly implemented the FD MSS program.




FIG. 2. Magnitude of pressure against distance from the stapes for five input frequencies. The pressures were normalized to be 1.0 at the stapes. Solid lines, FE MSS; dashed lines, FD MSS; frequencies indicated on graph. The figure shows strong qualitative similarity to Fig. 2 in Viergever (1977).

B. Dimensional accuracy

Having verified that FD MSS is able to produce the published results, we move on to show that FE MSS produces the accuracy improvement expected from a 2D model. We show that results from FE MSS show features typical of 2D models whereas FD MSS does not.

Neely (1981) described a 2D FD cochlear model in the frequency domain. Figure 3 in that paper shows a plot of the magnitude (in dB) and phase of the pressure versus distance from the stapes. Neely noted that in the 2D results “the magnitude of the pressure changes relatively little before the characteristic place and very abruptly just beyond that, followed by a break in slope to a more gradual decline in pressure (called the plateau region).” The start of this plateau region is most clearly seen in the phase data. These characteristics are also apparent in Fig. 1 for the FE MSS response

but the FD MSS response shows behavior typical of a 1D model. The data for this figure was produced by postprocessing the frequency response of the models from their final state space matrices. The magnitude is plotted on a dB scale $re 2 \times 10^{-5} \text{ Nm}^{-2}$ (rms), the eardrum pressure P_e . The frequencies used in this figure are indicated on the left hand side of the 2D results line in the phase plot. The same frequencies were used in the 1D case.

The FE MSS responses to 3 kHz and 750 Hz inputs show a clear “notch” after the best place. At some other frequencies, such as 1.5 kHz there is no such notch. A similar effect is visible in results from Neely (1981) and this behavior does not occur in 1D models (therefore the FD MSS code does not capture this feature). This feature has previously been discussed by de Boer and Viergever (1982, p. 146); Parthasarathi *et al.* (2000) and Watts (2000). de Boer and Viergever suggest micro-scale reflections might be occurring

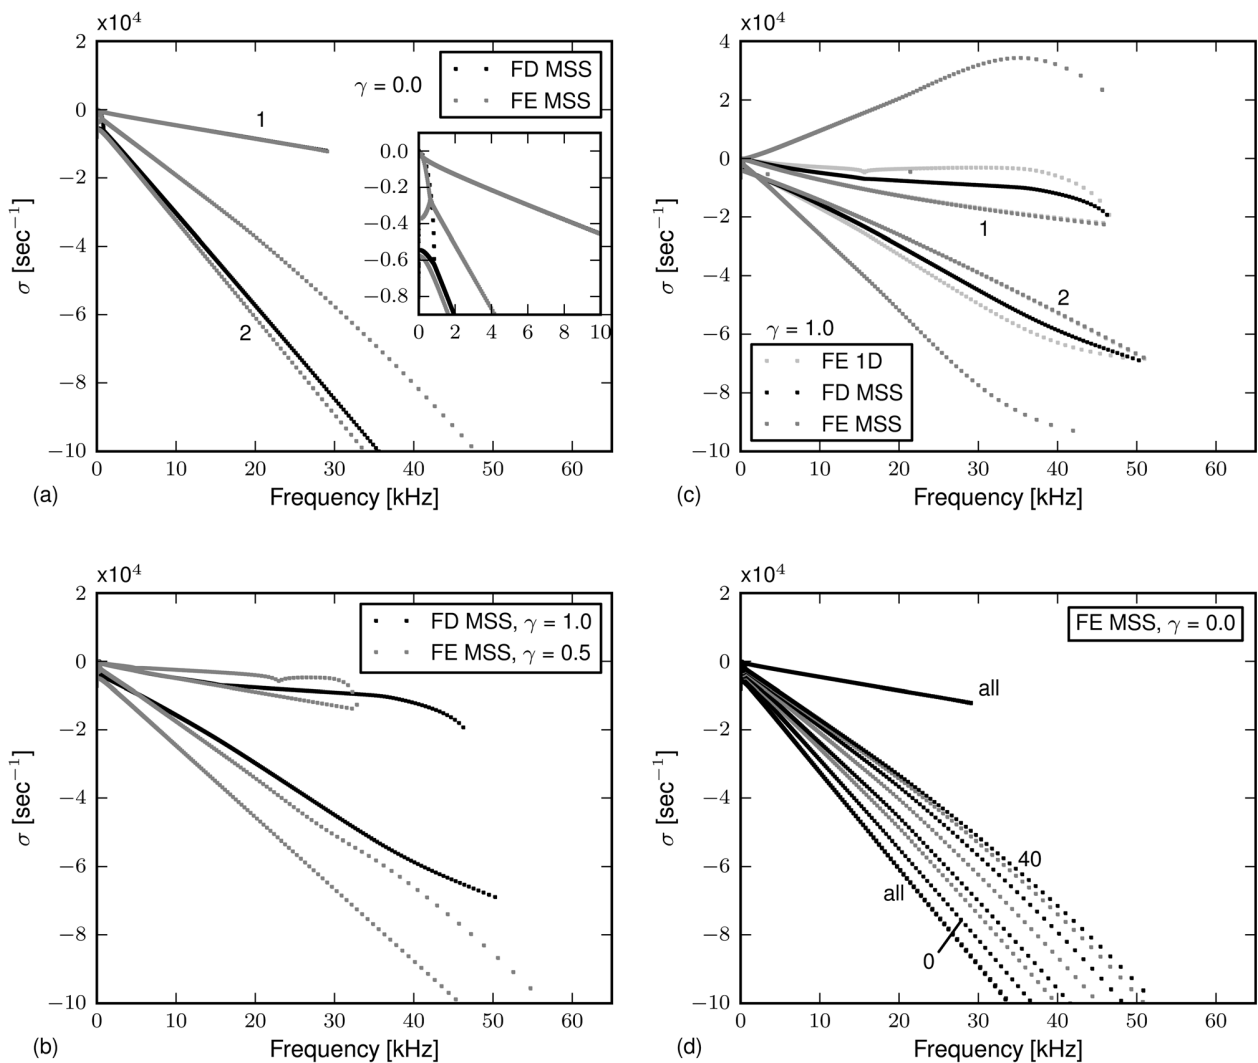


FIG. 3. Pole locations of the FD MSS and FE MSS programs for various choices of parameters. Each pole is a complex number $\sigma + j\omega$. The poles are plotted using the representation from Elliott *et al.* (2007) for convenient comparison to that paper. This representation puts the real part, σ , on the y axis of the figure and plots $f = \omega/2\pi$ on the x axis. The top two panels show (a) the passive case, $\gamma = 0$ in both codes, and (c) an active case, $\gamma = 1$ in all codes. The numbers “1” and “2” mark lines of poles that discussed further in the text. The inset in (a) shows more detail for the frequency range 0–10 kHz. The lower two panels show (b) another active case, $\gamma = 1$ for FD MSS and $\gamma = 0.5$ for FE MSS and (d) the poles of FE MSS for various numbers of cells in the vertical (H) direction. For (d) the model is passive ($\gamma = 0$) and the line of poles marked “0” is from a 1D FEM model. See the text for further description of the number of cells used.

due to the rapid change from short-wave to long-wave behavior. Parthasarathi *et al.* note that the notch appears in 3D models as well as in 2D models. They investigated the 4.53 kHz notch in Neely's 1981 model using a fine mesh (1000×50 FEM cells) and found that the notch was visible on a coarse mesh but did not occur on the fine mesh. They concluded that notches are the result of insufficient discretization. We investigated this further by rerunning the code for Fig. 1 using 1000×50 FEM cells as well (not shown). On the finer mesh we found notches occurring at 10 kHz and 800 Hz, but no evidence of a notch near 3 kHz. This suggests that Parthasarathi *et al.* might have missed a notch at a slightly different frequency or that the frequency of the notch might have shifted significantly. Watts suggests a possible explanation for the occurrence of notches in terms of destructive interference of wave modes in the cochlea, that is between the main traveling wave mode of the Liouville-Green approximation and the modes neglected by that approximation. Furthermore, Watts discusses literature that suggests similar behavior is seen in biological measurements. Therefore, the FE MSS program correctly captures the pertinent characteristics of a 2D cochlear model.

Viergever (1977) describes an investigation into the differences between 1D and 2D cochlear models published in the 1970s. Using a heuristic approach based on the Taylor expansion allowed Max Viergever to solve this early 2D model. He found large quantitative differences between the 1D and 2D case, although there were "hardly any" qualitative differences. Figure 2 in that paper shows plots of pressure magnitude versus distance along the BM for various input frequencies. We present a similar figure for data from FD MSS and FE MSS (Fig. 2). In this case the pressures were normalized so that the stapes pressure was 1.0. Some qualitative similarities to Viergever's results can be seen. At the lowest frequency (25 Hz) the 1D and 2D results are similar, but at higher frequencies there are quantitative differences. The 2D response for frequencies 10 kHz, 5 kHz, 3 kHz, and 500 Hz show a steeper drop-off in pressure in the vicinity of the best place than in Viergever (1977). This different steepness of response between the 1D and 2D results is consistent with the 2D model capturing the short-wave behavior of the model whereas the 1D model only captures the long-wave behavior (de Boer, 1991, Sec. 2.3). The two pressure profiles for a given frequency typically cross, which is only visible for the highest two frequencies in Viergever (1977). The responses to frequencies 25 Hz, 75 Hz and to a lesser extent 500 Hz show the *whole cochlear resonances* due to the apical boundary condition applied to the model. The model of the cochlear partition used in that report contains two states per element, whereas Neely and Kim (1986) used four states per element. Therefore, although there are qualitative and quantitative differences between the figures, the similarities support the fact that FE MSS is a fully 2D model.

In this section we have shown that the FE MSS program produces results with features that are characteristic of a 2D cochlear model. It is sufficient to compare just FE MSS and FD MSS because we show that the results of FE MSS are identical to FE MSSV to the expected numerical precision.

C. Effect of dimension and activity on model poles

The section above has shown that FE MSS is a fully 2D cochlear model. In this section we take advantage of the ability to plot the poles of the model to visualize the effect of dimension. The program FE MSSV could have been used instead of FE MSS to produce the pole plots, and would have produced identical results.

Figures 3(a) and 3(c) show a comparison between the FD MSS and FE MSS pole locations for the passive and active model, respectively. In the passive case (a) the dominant poles for both models lie along the same line in the figure (marked "1"). For FD MSS a second line of nondominant poles can be seen and this line continues off the figure. For FE MSS there are two lines, an upper and a lower (marked "2"). The poles of the isolated segments (that is, the poles of A_E) lie in close proximity to lines "1" and "2" and would overlap those lines if plotted on the same figure. The position of the poles in line "2" have not been visibly shifted by the fluid coupling in the FE MSS case.

The apparent similarity of the dominant poles may seem inconsistent with the results presented above showing a significant difference in their responses, particularly visible in Fig. 1. This is partly because the frequencies used for that figure lie in the range 0–10 kHz which is on the extreme left of Fig. 3(a). Therefore we include an inset that shows the pole locations for this frequency region. Once again the major difference between the two plots is the location of the nondominant poles. This is a reminder that the poles of a system alone do not determine the frequency response because the system's "zeros" also have a significant effect. The behavior of systems with a couple of poles and zeros can be reliably predicted "by eye" from their locations; however, this does not extend to such large systems. Therefore, these plots do not provide enough information to make frequency response tests unnecessary. Nevertheless the poles of a system have a number of uses: they immediately show whether a model is stable, and they were examined to determine the frequency of the *whole cochlear resonance* when plotting Fig. 2.

In the active case [Fig. 3(c)], the FD MSS model is stable because for all poles the real part σ is negative. The FE MSS model is highly unstable because there are a large number of poles with σ positive and the maximum value of σ is of the order of 3×10^4 . This value for σ implies that the envelope of the response of the system becomes unbounded as $\approx e^{3 \times 10^4 t}$, which agrees well with time domain results for this system. Once again the lines marked "1" and "2" indicate lines of poles from the FE MSS model that are not visibly shifted from poles of A_E by the fluid. The variation between these two models could have two possible causes: both the dimension of the model and the fluid discretization method (FD or FEM) can affect the pole locations. In order to show the relative effect of these two factors we include the poles from a 1D variant of FE MSS (FE 1D). Two lines of poles from the FE 1D model are visible in the figure and another two lines lie beneath the poles of the FE MSS model marked "1" and "2." Although the values of σ for the dominant poles in the FE 1D model are closer to zero than for the FD MSS model, they remain negative and hence stable.

Figure 3(b) shows the effect of reducing γ to 0.5 in the 2D case (this is still an active case). The poles for $\gamma = 1$ in the 1D case are plotted for comparison. This value of feedback gain gives approximately the same stability as the active 1D case, but the range of frequencies covered by the dominant poles is significantly reduced. Therefore, the model would have a different frequency response. This shows that tuning γ in the 2D model is not sufficient to match the dominant pole locations of the 1D active case. It is also a reminder that parameters fitted to a 1D model can produce unexpected results when extended beyond their operating range.

Finally, Fig. 3(d) shows the effect of varying the number of cells in the height H direction. In all cases some of the poles lie along two lines marked “all.” These are the same lines as “1” and “2” in Fig. 3(a). The remaining poles lie on another line that changes position depending on the number of cells in the vertical direction. We present results for 1, 2, 5, 10, 20, and 40 cells. The pole locations of the FE 1D model are also included in the figure and are marked by “0.” The remaining lines lie between these lines in sequence from the line marked “0,” with the final line of poles indicated by “40.” As the number of cells in the vertical direction is reduced, the second line of nondominant poles approaches the FE 1D poles marked. The nondominant poles of the FD MSS case lie between the lines of nondominant poles marked “all” and “0.” Therefore, as the number of cells in the vertical direction of the 2D model are reduced the pole locations approach the 1D case; however, there is still a significant difference between FE 1D case and FE MSS with one cell in the vertical direction. Some qualitative differences are to be expected because the 1D and 2D models make fundamentally different assumptions about the vertical acceleration of the fluid. In the 1D model only horizontal acceleration of the fluid is allowed, leading to a paradox at the cochlear partition boundary (see de Boer, 1980, Sec. 6.2), whereas in the 2D model vertical acceleration of the fluid is allowed, except along the top boundary. Ten FEM cells in the height direction was the selected value for the remainder of the results because this makes for roughly square cells.

In these figures we have seen that different fluid discretization strategies shift the poles of \mathbf{A}_E in different ways. In FD MSS *all* the poles are shifted by a comparatively small distance. When FEM is used to discretize the fluid some poles are shifted by a large distance and others relatively little. The distance that poles are shifted by depends on whether a 1D or 2D model is used, and as the number of cells is reduced in a 2D model the pole positions tend towards the 1D poles. The absolute shift in poles between the passive and active case, Figs. 3(a) and 3(c), is similar for a given fluid discretization, however, the active case is more sensitive to this shift because it is so much closer to instability.

D. Computational performance of MSS and MSSV methods

1. Numerical precision

In order to assess the numerical precision of these methods, we consider sources of precision loss in the processes

required to form each \mathbf{A} matrix. There are two major sources that are relevant. Firstly, when matrix equations such as $\mathbf{A}\mathbf{x} = \mathbf{b}$ are solved for \mathbf{x} there is a precision loss that depends on the condition of the \mathbf{A} matrix (Cheney and Kincaid, 2008, p. 321). Secondly, when two variables that are identical to the n th significant figure are subtracted the precision of the result is reduced by n significant figures (Cheney and Kincaid, 2008, Sec. 2.2). In the MSSV method just one comparatively well conditioned matrix needs to be inverted rather than two comparatively poorly conditioned matrices for the MSS method. This suggests that MSSV is more numerically accurate, particularly for large systems. Loss of significance due to subtraction was not a problem for either method for this model. Appendix B provides details of how these results were obtained.

To test whether the methods are identical to numerical precision we resort to a smaller problem of 4×2 cells, because it is difficult to bound the number of significant figures lost in each algorithm. We calculate the \mathbf{A} matrices in double precision using each algorithm and then compare their single precision representation. The \mathbf{A} matrices computed using the MSSV and MSS methods were found to be identical to single precision for the smaller problem, suggesting that the methods are identical to numerical precision. Once again Appendix B provides further details.

2. Storage and computational requirements

The methods FE MSS, FE MSSV and FE MI place different requirements on the machine used to produce the results. Appendix C provides a theoretical assessment of their requirements and how it is affected by fluid dimension and mesh refinement. For storage requirements, the assessment focuses on the size of the largest data structures: $\mathbf{K} - \mathbf{R}\mathbf{C}_E\mathbf{B}_E\mathbf{Q}$ in the case of FE MI and \mathbf{A} for FE MSS and FE MSSV. See Table III in that section for a summary. The computation effort for each method is dominated by the matrix inversions required. We focus on LU factorization for the assessment but also consider the conjugate gradient method where appropriate. A set of “worst-case” parameters from a large cochlear model is used to provide firm numbers to compare the methods.

It is found that, while the asymptotic order of operations often favors the FE MSS and FE MSSV codes, the figures from the “worst-case” calculation generally favor FE MI code, the exception being for large 3D cochlear models. LU factorization is feasible for even the largest 2D systems considered but the conjugate gradient method would be required for large 3D systems. The major distinction between the FE MSS and FE MSSV is the process used to assemble \mathbf{A} . Once again, the fact that FE MSS requires an extra matrix inversion (which is full) means that the FE MSSV method would be the preferred method to assemble \mathbf{A} . Table IV summarizes our findings.

E. Time domain results

Finally, we present time domain results from the programs. The major point we wish to highlight in this section is that the FE MI program, which never forms the complete

state space matrices, produces the same time domain results as FE MSS and FE MSSV.

Figure 4 shows the response of the pressure at the cochlear partition, location $x = 5.02$ mm, for the input:

$$q_1 = \begin{cases} 1 \text{ ms}^{-2}, & 0 \leq t < 1 \mu\text{s} \\ 0 & \text{otherwise.} \end{cases} \quad (59)$$

This input approximates a weighted impulse function, or a click in physiological experiments, and contains a wide band of frequencies. The system is linear so the response can be rescaled as desired for comparison with other results. Figs. 4a and 4b shows the behavior of the models over a short time period (1 ms) and relatively long simulation (0.1 s), respectively. In all cases the responses for FE MSS, FE MSSV and FE MI overlap on the plot and are only visible as a single line, showing that they produce equivalent results in the time domain. (This is consistent with the findings of Sec. VB because only a few significant figures can be reliably distinguished in a plot.)

In Fig. 4(a) the response of FD MSS is clearly distinct from the remaining plots. Results from the 1D variant of FE MSS produces agree with the FD MSS response so the difference is due to the fluid dimension primarily. In a plot of the response over a short time period the high frequency response of the models dominates (in particular the characteristic frequency of the element). Over the longer time scales shown in Fig. 4(b) the differences between the codes are less obvious, but the inset shows that they are still present. These results correspond well to Fig. 2 which showed large differences due to dimension at high frequencies but much smaller differences at low frequencies.

In Fig. 4(b) the low frequency *whole cochlear resonance* dominates the response. It is important for a time domain cochlear model to perform well in the face of whole cochlear resonance (and these models do because they all agree concerning its dynamics). However, the resonance for this data set is larger than would be expected in the cochlea. This is because the data set originates from a frequency domain model in Neely and Kim (1986) where apical reflections and whole cochlear resonance is not problematic. Therefore, simple precautions that can minimize apical reflections [such as those discussed in Puria and Allen (1991)] were not used.

Figure 5 shows cochlear partition acceleration for the same test conditions. In this case the spatial response is shown for a sequence of time instants. The envelope of the response shifts apically for increasing time as expected. The envelope of the 1D FD MSS response (shown in gray) is slightly apical to the envelope of the 2D codes. (Once again the FE 1D code produces similar results to the FD MSS code.) The difference between the 1D and 2D response is consistent with Steele and Taber (1979, Fig. 9), which shows that for a given frequency the envelope of the response for a 1D model is shifted apically compared to that of a 2D model. At 2 ms the envelope has reached the apex of the cochlea and the apical boundary condition is beginning to affect the shape of the wave form. This is more pronounced at 4 ms and we do not show any time instants beyond 4 ms because

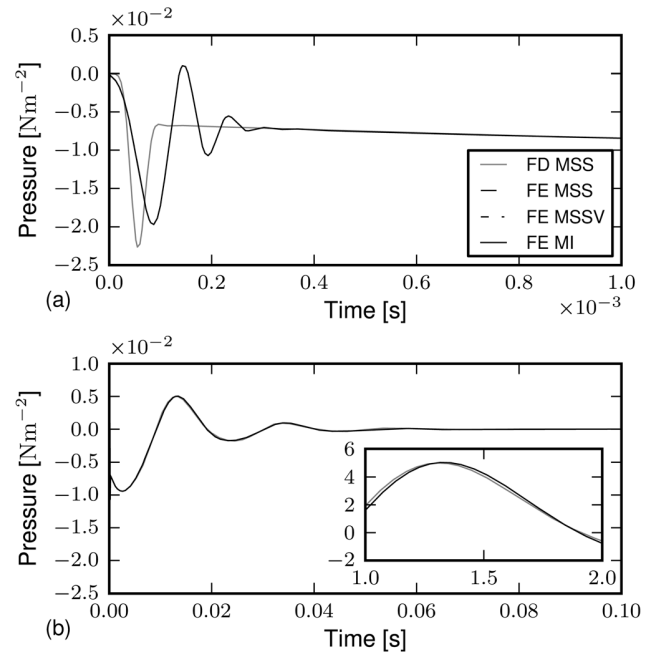


FIG. 4. Response of the pressure at point $x = 5.02$ mm on the cochlear partition to a unit (1 ms^{-2}) $1 \mu\text{s}$ pulse occurring at $t = 0$. The feedback gain is set as $\gamma = 0$ and zero pressure is enforced at the helicotrema. Note that the response lines for FE MSS, FE MSSV, and FE MI overlap and are only visible as a single line in the figure. (a) The initial response for the first 1 ms of the simulation. (b) The response over a longer time period: 0.1 s. The inset, which shows the period 10 to 20 ms, highlights the fact that the FD MSS code still produces slightly different results on this time scale.

they are dominated by the apical reflections and are no longer representative of expected cochlear behavior or time domain models in general.

For the results presented in Figs. 4 and 5, zero pressure was enforced at the helicotrema. The FE MSSV and FE MI codes ran to completion and gave meaningful results even if a zero normal flow boundary condition was enforced at the helicotrema, whereas the MSS codes fail for this boundary condition. The MSS codes fail because there are insufficient boundary conditions imposed to invert \mathbf{F} so the programs that require \mathbf{F}^{-1} throw numerical exceptions. FE MSSV and FE MI never form \mathbf{F}^{-1} because it is not required for the algorithm in Diependaal *et al.* (1987) so these programs can run to completion.

VI. DISCUSSION

A. Dimension and geometry

As noted in the introduction, there is support in the literature for the improved accuracy of 2D models over 1D models. The results in Sec. VB show characteristic 2D features for the 2D model and the expected differences between the 1D and 2D models. Time domain results in Sec. VE show comparable dimensional effects. Therefore, the proposed extension captures the behavior of a 2D fluid. A similar method can also be used to extend Elliott *et al.* (2007) to 3D models; therefore, the discussion of computational performance also considers the 3D case.




FIG. 5. Cochlear partition acceleration at five time instants for the input given by (59). Units of the acceleration are ms^{-2} . Black line: response of FE MSS model (this is representative of FE MSSV and FE MI models as Fig. 4 shows). Gray line: response of FD MSS model.

In the extension to 2D, additional variables describing the behavior of the oval window are introduced (located at quadrature points of the finite element grid). These can be related to each other by a function describing the movement of the oval window as a whole so that its behavior can be summarized by one number as before. Similarly, in the extension to a 3D fluid model some assumption about the deflection of the basilar membrane in the radial dimension can be made (e.g., see Parthasarathi *et al.*, 2000; Lim and Steele, 2002).

Neely and Kim (1986) note that their model can be generalized to 2D or 3D using FD. This suggests that a FD extension of Elliott *et al.* (2007) to 2D models is also feasible. The finite element method offers greater freedom of geometry than methods that use finite differences or Green's functions. Therefore, the FEM based \mathbf{F}^{-1} matrices are more flexible than FD equivalents. There is also evidence that the fluid geometry has an effect on accuracy (Puria and Allen, 1991; Cai *et al.*, 2005) so it is expected that this flexibility will be useful.

The size of the \mathbf{F}^{-1} matrix is unchanged when using the 2D FEM approach (a few additional variables are assigned on the oval window but they can be eliminated as described above). This lends a certain elegance to the FEM extension since the theory developed for the FD case can continue to be used. The proposed \mathbf{F}^{-1} matrices require more computational operations to assemble, but this is offset by the fact that they can be used multiple times.

While 2D FEM based \mathbf{F}^{-1} matrices can simply replace 1D equivalents, the dimensional change does have an effect on cochlear tuning beyond the desired accuracy improvements. This fact is clear from the computational results presented, where shifts in the best place for a frequency and even instability in a previously stable model have been shown. Therefore, while minimal changes are required to the

programs used, 1D tuning parameters do need to be adjusted for use in 2D models.

B. MSS, MSSV, and MI

We have claimed that the MI method, which we used Diependaal *et al.* (1987) to outline, is representative of a large number of other time domain cochlear models. To support this claim we look for equations that group terms in a similar manner to Eq. (44) above. Terms similar to either the left or right hand side of this equation are generally evidence that a similar approach was taken. We find that Allen and Sondhi [1979, Eqs. (27a) and (27b)], Matthews [1980, Eqs. (2.7.a), (8.21)], Diependaal and Viergever [1989, Eqs. (10), (11)], van Hengel [1996, p. 76, Eq. (7)], and Liu and Neely [2010, Eq. (18)] match this criteria. Given the extent that the method introduced by Allen and Sondhi (1979) has been adopted by other authors, we believe it is fair to say it was the only method of its type in use before Elliott *et al.* (2007) for the class of models it applies to. We are not aware of any earlier alternatives.

In Sec. IV we show that the MI method can also be expressed in terms of the state space matrices used for the MSS method. This led to a variation of the MSS method that we called MSSV. We show that MSS and MSSV are related by an identity from Henderson and Searle (1981). The identity presented suggests that MSS and MSSV are identical to the limit of numerical precision. This is supported by numerical results. The MI method itself is useful for the time domain only, but it produces identical results to MSSV without requiring large state space matrices to be formed and stored. An analysis of the computational performance of the various methods shows that the MI method offers the best performance in all respects for 2D models, but the MSSV and MSS methods become competitive for large 3D models.

Both MSS and MSSV form the \mathbf{A} matrix explicitly. This is useful for determining the stability of the linear system from its eigenvalues. Recent papers leverage this (Ku *et al.*, 2009; How *et al.*, 2010). Similarly, once the model is in the final state space form either time domain or frequency domain results can be computed. The drawback of forming \mathbf{A} explicitly is that it is full and must be stored. The analysis of the codes confirmed that the storage requirements of \mathbf{A} are much larger than for \mathbf{K} or $\mathbf{K} - \mathbf{R}_E \mathbf{B}_E \mathbf{Q}$, except for large 3D models (see also Bertaccini and Sisto, 2011).

We now ask what differences between MSS and MSSV users should be aware of. The defining difference between the two methods is the fact that MSSV only requires a single inversion of a sparse matrix, whereas MSS requires an extra inversion of a full matrix. This is a direct result of the identity relating the two methods, Eq. (51). The second inversion required means that using MSS to assemble \mathbf{A} is more computationally intensive than MSSV. Both the MSS and MSSV methods are sensitive to precision loss during computation and preconditioners can be used to improve the condition numbers of the matrices in all of the methods. The method proposed by Elliott *et al.* (2007) is more sensitive to loss of precision than the variation proposed in this paper due to the extra inversion, as well as the relative condition of \mathbf{K} and

$\mathbf{K} - \mathbf{R}\mathbf{C}_E\mathbf{B}_E\mathbf{Q}$. Although these tests were only performed for the model and parameters studied, we expect the result to extend to other models because the matrices considered are characteristic of a wide range of cochlear models. The term $-\mathbf{R}\mathbf{C}_E\mathbf{B}_E\mathbf{Q}$ that appears in the MSSV method [see Eq. (57)] forms part of a *Robin* boundary condition along the cochlear partition. Therefore, the MSSV method makes use of the information available about the cochlear partition boundary at an earlier stage in the computation than the MSS method which forms \mathbf{F}^{-1} directly. This is part of the reason that the matrices in the MSSV method are better conditioned than in the MSS method. The fact that the apical boundary condition is not required for numerical stability of the MSSV method is further evidence that the cochlear partition boundary is having a stabilizing effect. We therefore suggest that MSSV is a superior method of assembling the state space matrices. An additional feature of the MSSV method is that it allows the stapes and helicotrema boundaries to be treated in the same manner as the cochlear partition boundary.

Having noted these differences, the Eqs. (57) and (58) allow us to move between the MSS and MSSV methods so we can select the expression most appropriate for our needs when analyzing cochlear models. If the eigenvalues are required in an MI style computation, Eq. (57) shows how the required \mathbf{A} matrix can be assembled. Should we wish to use the MSS method in a case where \mathbf{K} is not invertible, we can assemble the \mathbf{A} matrix via the MSSV approach and continue as normal. Finally, the identity suggests how we can plan the work in the state space formalism (using either MSS or MSSV) and then work back to an MI implementation to minimize storage and computational requirements.

C. Recent developments

Sisto *et al.* (2010) gives an overview of how a wide range of current nonlinear cochlear models can be implemented in the framework of Elliott *et al.* (2007). Bertaccini and Sisto (2011) describes their method in more detail and in particular focuses on the question of efficient computation. These papers represent a significant contribution to the cochlear modeling framework. Their approach offers full second order convergence properties (in the spatial domain), as well as a more efficient solver based on Krylov space methods [see the motivation for these methods in Bertaccini and Sisto (2011)]. We note that the FEM approach presented in this paper also offers full second order convergence properties (provided that the shape functions used are at least linear). Krylov space methods have been used in cochlear models previously in the form of the conjugate gradient method, for example Kolston and Ashmore (1996). It can be used as the spatial solver in models following Diependaal *et al.* (1987). Appendix C shows that the conjugate gradient method can be applied to any of the methods discussed in this paper and is useful for 3D problems.

Finally Sisto *et al.* (2010) state that the implicit time stepping algorithm they propose is as efficient as explicit schemes because the mass matrix they form is nontrivial and

explicit schemes would require the solution to a system of nonlinear equations at each time step. Their paper addresses a wide range of cochlear models, and we wish to clarify that the implicit time stepping algorithm is only superior for some of these models. Firstly, we note that a number of nonlinear cochlear models have been solved with an explicit time stepping scheme. These include Duifhuis *et al.* (1986), Diependaal and Viergever (1989), Liu and Neely (2010), as well as the nonlinear example in Elliott *et al.* (2007) itself. None of these papers mention any difficulties with solving the nonlinear models using an explicit scheme. To examine why implicit methods were not necessary for these models we refer to Eq. (57). Note that both ways of expressing the \mathbf{A} matrix is of the form of an expression dependent on \mathbf{B}_E , \mathbf{C}_E and the fluid matrices, multiplied by the matrix \mathbf{A}_E . There are matrix inversions in the first part of the expression but \mathbf{A}_E itself is never inverted. The nonlinear models mentioned above only include nonlinear terms that affect the \mathbf{A}_E matrix. These terms are effectively nonlinear stiffnesses and dampings in the cases mentioned. The nonlinearities affecting \mathbf{A}_E are computationally cheap to deal with in an explicit method because it is never inverted. Therefore, any nonlinear “non-local”⁴ coupling that can be restricted to the \mathbf{A}_E matrix might not require an implicit time stepping scheme. On the other hand, \mathbf{B}_E appears in expressions that must be inverted in both forms of the \mathbf{A} matrix. If any nonlinear terms affect \mathbf{B}_E , these inverses would need to be computed multiple times when advancing a single time step and the comments in Bertaccini and Sisto (2011) regarding stiffness and implicit time stepping apply. \mathbf{B}_E includes the mass of the element and terms which are mass-like in the sense that their contribution to the element’s acceleration is dependent on the pressure acting on the element. The literature on feed forward mechanisms (Geisler and Sang, 1995; Lim and Steele, 2002; Kim and Xin, 2005) formulates a nonlinear function that has a mass-like characteristic as well as damping and stiffness-like characteristics. This is one of the types of models addressed in Sisto *et al.* (2010).

Therefore, the implicit method proposed by Sisto *et al.* (2010) is not always required, but we are pleased to see a useful extension to Elliott *et al.* (2007) for systems where it is. Equation (57) can be used to distinguish when the implicit method is required and when an explicit method is likely to be adequate.

VII. CONCLUSIONS

We develop FEM matrices to replace the 1D FD fluid model \mathbf{F}^{-1} in Elliott *et al.* (2007) with a 2D model. The replacements offer full second order convergence properties and greater freedom with the shape of the fluid region, due to the properties of FEM. The literature suggests that using a 2D fluid model will result in an automatic increase in physiological accuracy and our numerical results show the expected difference between the 1D and 2D models. Once the 2D matrices are computed, they may be used multiple times. The relationship between the methods Elliott *et al.* (2007) and Diependaal *et al.* (1987) has been discussed. It is shown that these methods are identical in some senses and

complementary in others. We derive an alternative method to construct the required state space matrices that appears to offer improved numerical accuracy and performance. Equations relating the methods are given here that allow models to be converted from one form to the other in order to utilize their specific strengths. We note that the implicit method proposed by Bertaccini and Sisto (2011) is not required for certain classes of nonlinear cochlear models but does allow new classes to be tackled.

ACKNOWLEDGMENTS

One of us (M.J.R.) would like to thank the deal.II users group in CERECAM for the many fruitful discussions on topics relevant to this paper. The financial assistance of the National Research Foundation (NRF, South Africa) towards this research is hereby acknowledged. Opinions and conclusions arrived at are those of the authors and are not necessarily to be attributed to the NRF.

APPENDIX A: INTEGRATION OF STAPES PRESSURE

In Sec. III the FEM matrix \mathbf{Q}_s ($n_s \times n$) was described. Its function is to interpolate a vector of pressures at the FEM nodes to the quadrature points. The user is then free to decide how to use this data. For example n_s state space models could be used to define the behavior of this boundary, using the data directly. The stapes is often modeled as a single moving unit. In this case it is the integral of the pressure that causes the force. Therefore,

$$f_s(t) = \int_{\Gamma_s} p(x, y, t) d\Gamma_s, \quad (\text{A1})$$

where $f_s(t)$ is the force acting on the stapes. Some models convert from pressure to force in the stapes state space model. In these cases the average pressure can be obtained by dividing $f_s(t)$ by the stapes area. The programs presented in Sec. V convert from force to pressure by dividing through by the cochlear height.

We use the FEM shape function Eq. (11) and Gauss quadrature to obtain

$$\begin{aligned} f_s(t) &= \int_{\Gamma_s} \mathbf{N}(x, y) d\Gamma_s \mathbf{p}_n(t) \\ &= \sum_{i=1}^{n_s} W_i |\mathbf{J}(\xi_i)| \mathbf{N}(\xi_i) \mathbf{p}_n(t). \end{aligned} \quad (\text{A2})$$

This shows that

$$\bar{\mathbf{Q}}_s = \sum_{i=1}^{n_s} W_i |\mathbf{J}(\xi_i)| \mathbf{N}(\xi_i), \quad (\text{A3})$$

which is $1 \times n$. Alternatively, $\bar{\mathbf{Q}}_s$ may be manipulated to include \mathbf{Q}_s as follows:

$$\begin{aligned} f_s(t) &= \sum_{i=1}^{n_s} W_i |\mathbf{J}(\xi_i)| \mathbf{N}(\xi_i) \mathbf{p}_n(t) \\ &= \sum_{i=1}^{n_s} W_i |\mathbf{J}(\xi_i)| \mathbf{IN}(\xi_i) \mathbf{p}_n(t) \\ &= \mathfrak{W}_s \mathbf{Q}_s \mathbf{p}_n(t), \end{aligned} \quad (\text{A4})$$

where $\mathfrak{W}_s = W_i |\mathbf{J}(x_i)| \mathbf{I}$ is $1 \times n_s$ because \mathbf{Q}_s is $n_s \times n$. Then $\bar{\mathbf{Q}}_s = \mathfrak{W}_s \mathbf{Q}_s$.

In practice it is convenient to make use of the similarity between $\bar{\mathbf{Q}}_s$ and $\bar{\mathbf{R}}_s$. Note that in this case $\bar{\mathbf{Q}}_s = -\bar{\mathbf{R}}_s^T / 2\rho$ because the same shape functions are used for the weighting functions and trial solutions.

APPENDIX B: NUMERICAL PRECISION

In this section we consider whether the MSS or MSSV method provides superior numerical properties for computing the required \mathbf{A} matrix. We consider the condition number of matrices that require inversion, and potential sources of loss of significance due to subtraction. Finally, we discuss whether the methods agree to numerical precision.

1. Matrix condition number

Consider the general matrix equation $\mathbf{Ax} = \mathbf{b}$. The condition number of a matrix κ provides a measure of the effect of errors in \mathbf{b} on the solution \mathbf{x} . The ideal condition number κ for a matrix is 1.0 and a general rule is that an inversion loses one digit of precision for each order of magnitude increase in the condition number (Cheney and Kincaid, 2008, p. 321).⁵ FE MSS requires the inverse of \mathbf{K} and $\mathbf{I} - \mathbf{B}_E \mathbf{Q} \mathbf{K}^{-1} \mathbf{R} \mathbf{C}_E$. FE MSSV requires only one inversion, for the matrix $\mathbf{K} - \mathbf{R} \mathbf{C}_E \mathbf{B}_E \mathbf{Q}$. The condition numbers of these matrices are shown in Table II for the full simulation (249×10 cells) and a smaller simulation (4×2 cells). The final entry in the table will be discussed in Sec. B.3 below. The condition numbers of the matrices required for the FE MSS program are both at least an order of magnitude larger than the condition number of the matrix required for the FE MSSV program. Therefore, based on the condition numbers, the MSSV method is a more accurate way to compute the \mathbf{A} matrix.

2. Loss of significance due to subtraction

Loss of significance can occur when values are subtracted, or values with opposite sign are added (Cheney and Kincaid, 2008, Sec. 2.2). We check whether loss of significance occurs during these processes for the 249×10 cells case. Three potential locations were checked: $\mathbf{K} - \mathbf{R} \mathbf{C}_E \mathbf{B}_E \mathbf{Q}$ and $\mathbf{I} + \mathbf{B}_E \mathbf{Q} [\mathbf{K} - \mathbf{R} \mathbf{C}_E \mathbf{B}_E \mathbf{Q}]^{-1} \mathbf{R} \mathbf{C}_E$ in FE MSSV, and $\mathbf{I} - \mathbf{B}_E \mathbf{Q} \mathbf{K}^{-1} \mathbf{R} \mathbf{C}_E$ in FE MSS. We first checked the minimum absolute value created by each of the subtractions. This gives a measure of whether loss of significance is occurring for the particular model size. The minimum values obtained for the operations required for FE MSSV were 0.216 and 0.334, respectively. Neither of these values suggest that loss of significance is occurring. In the FE MSS code, the minimum value was 6.76×10^{-3} , which suggests that some loss of significance may be occurring.

To follow up on this finding we observed that changes in either the fluid grid and geometry or the values in \mathbf{B}_E and \mathbf{C}_E could affect this minimum value. For the model studied, the width of the basilar membrane appears in the expression for \mathbf{C}_E and by decreasing its value by $\approx 0.67\%$ we were able to obtain a minimum absolute value on the order of 1×10^{-16} . Even in

this worst-case scenario $\mathbf{I} - \mathbf{B}_E \mathbf{Q} \mathbf{K}^{-1} \mathbf{R} \mathbf{C}_E$ was invertible and the numerical results were meaningful. Therefore, this potential loss of significance does not have a serious effect on the MSS method.

This shows that loss of significance due to subtraction is not a concern for either the MSS method or the MSSV method. The matrices discussed contain terms that are shared by a wide range of cochlear models so these results are likely to extend to other models.

3. Equivalence of MSS and MSSV methods

Finally, we discuss the simple test to determine whether the MSS and MSSV are equivalent to the limit of numerical precision. The major difficulty in asserting that they are is that the precision loss must be bounded. We will not attempt to quantify this bound formally. Instead, we note that for a model with 249×10 cells the \mathbf{A} matrices computed are not identical when represented in single precision. (See Table II.) This suggests that more than 10 significant figures (on the processor used to produce the results) have been lost in one (or both) of the methods for a problem of this size. The large condition numbers of all the matrices involved suggest that this is reasonable to expect. For a smaller problem of 4×2 cells the \mathbf{A} matrices are identical to single precision and the condition numbers would also suggest that less than 10 significant figures should have been lost. The mathematical identity (57) implies that the methods should be identical if floating point arithmetic was not required. This informal analysis supports the claim that the MSS and MSSV methods are identical to numerical precision.

APPENDIX C: COMPUTATIONAL PERFORMANCE

In this section we consider how the memory and computation time required for each method scales with fluid dimension and mesh refinement. Equations (57) and (58) summarize the MSS and MSSV methods. The MI method computes a new estimate for the acceleration of the cochlear partition using a sequence of three equations: (43), (44), (47). [Removing the multiplication by \mathbf{C}_E in (47) provides the state gradient for the system.] This analysis assumes that FEM with linear shape functions is used for the spatial discretization but the general trends are relevant to FD and other shape functions as well. For simplicity, a rectangular fluid domain is considered.

TABLE II. Table showing the condition number of various matrices and the actual precision loss in the \mathbf{A} matrices for the FE MSS and FE MSSV programs. The first two matrices are required for the FE MSS program and the third is required for FE MSSV. The final entry records whether the two \mathbf{A} matrices were identical when converted to single precision. The results are presented for the full simulation used elsewhere (249×10 cells) and a smaller simulation (4×2 cells).

Expression	249×10 cells	4×2 cells
$\kappa(\mathbf{K})$	1.10×10^5	1.34×10^3
$\kappa(\mathbf{I} - \mathbf{B}_E \mathbf{Q} \mathbf{K}^{-1} \mathbf{R} \mathbf{C}_E)$	1.99×10^4	3.39×10^3
$\kappa(\mathbf{K} - \mathbf{R} \mathbf{C}_E \mathbf{B}_E \mathbf{Q})$	7.15×10^2	1.29×10^2
\mathbf{A} (Single Precision)	Differ	Identical

We define some variables to characterize the fluid grid and cochlear partition elements. Once again let n be the number of FEM nodes in the domain and n_{cp} be the number of quadrature points on the cochlear partition. n_{cp} must be even in this case. We define the number of nodes in length, height and width dimensions as n_x , n_y , and n_z , respectively. Finally, let n_{states} be the number of states per cochlear partition element. For linear shape functions $n_x = n_{cp}/2 + 1$. n depends on dimension, it is n_x , $n_x n_y$ or $n_x n_y n_z$ in 1D, 2D or 3D, respectively. If we assume that $n_{cp} \gg 1$ and FEM cells are roughly square then $n_x \approx n_{cp}/2$, $n_y \approx \frac{H}{2L} n_{cp}$ and $n_z \approx \frac{W}{2L} n_{cp}$, where L , H , and W are the length, height and width of the cochlear model. Therefore, $n \propto n_{cp}^D$, where D is the dimension of the model.

We will substitute values for these variables to estimate the worst case memory and number of operations required for each method. Where required, we will use $n_{cp} = 3000$, $L = 25$ mm, $H = 1$ mm, $W = 1$ mm and $n_{states} = 4$. L , H , and W are appropriate for a cat cochlea, the general trends for other species are expected to be similar. When $n_{cp} \approx 3000$ there is a one to one correspondence between cochlear partition elements and inner hair cells. If n_{cp} is increased beyond this the physical interpretation of a cochlear partition element is lost. Therefore, we consider 3000 to be a natural upper limit on the number of elements a cochlear model is likely to have.

1. Basic storage requirements

For the MI method, the matrix $\mathbf{K} - \mathbf{R} \mathbf{C}_E \mathbf{B}_E \mathbf{Q}$ is still explicitly formed. It is an $n \times n$ sparse matrix but each node couples with at most 3^D other nodes in the mesh for the assumptions used. Therefore, it can be stored with $n \times 3^D$ double precision variables. (This neglects overhead associated with indexing the nonzero elements.) For the example above, it requires 8.1×10^5 and 1.458×10^8 double precision values in 2D and 3D, respectively. Its size is order n_{cp}^D , which we will write as $O(n_{cp}^D)$, due to its dependence on n . The storage required to store the state space matrices associated with the n_{cp} elements is dominated by the \mathbf{A}_n matrices that require $n_{cp} \times n_{states} \times n_{states}$. This is $O(n_{cp})$ and will generally be negligible.

Both MSS and MSSV produce the matrix \mathbf{A} that must be stored in RAM. It is a full matrix in general, approximately $n_{cp} n_{states} \times n_{cp} n_{states}$, where we have neglected states associated with the stapes boundary condition. Depending on the details of the cochlear partition state space equations, many of the terms in this matrix may be zero. For the 1994 state model discussed in Sec. V approximately 3/4 terms (2978041 of the 3976036 terms) are zero. However, standard linear algebra libraries do not provide convenient methods to avoid storing and performing computations with these zero terms. Efficient programming would store $\mathbf{B} \mathbf{F}^{-1}$ as a single matrix rather than separately. Although the combined matrix is approximately $n_{cp} n_{states} \times n_{cp}$, if $\mathbf{q}(t)$ has only one nonzero term (as is the case in the example) only $n_{cp} n_{states} \times 1$ needs to be stored. \mathbf{F}^{-1} , which is $n_{cp} \times n_{cp}$, is required to compute the pressure in the model. The size of these matrices are all independent of dimension. \mathbf{A} , which uses approximately

1.44×10^8 double precision numbers, requires the lion's share of the memory. Its memory requirements grow as $O(n_{cp}^2)$. While forming \mathbf{A} , $\mathbf{B}\mathbf{F}^{-1}$ and \mathbf{F}^{-1} we need to perform operations on either \mathbf{K} or $\mathbf{K} - \mathbf{R}\mathbf{C}_E\mathbf{B}_E\mathbf{Q}$. These matrices share the same sparsity pattern (locations of nonzero terms) so the analysis above is also applicable to \mathbf{K} .

Table III provides a summary of the largest data structures required for each method. The MSS and MSSV methods have similar storage requirements according to this first analysis. (We have neglected memory required RAM required by the numerical methods implementing the matrix inversions, which will be discussed below.) Their final memory requirements are dominated by the \mathbf{A} matrix. The MI method's storage requirements are dominated by the $\mathbf{K} - \mathbf{R}\mathbf{C}_E\mathbf{B}_E\mathbf{Q}$ matrix. The MSS and MSSV methods must also store a matrix of this size while constructing \mathbf{A} . Although the sizes of \mathbf{A} and \mathbf{K} are $O(n_{cp}^2)$ and $O(n_{cp}^D)$, respectively, consideration of the worst case values selected shows that \mathbf{A} is larger than \mathbf{K} for all but large 3D models. Therefore, the MSS and MSSV will generally require more storage for simply representing data than the MI method.

2. Computational effort

The work required for a time domain simulation consists of a once off initialization procedure followed by repeated computations of the next solution. In this section we neglect any postprocessing stages as they are similar for all methods. If the Runge-Kutta 4th (RK4) order algorithm is used to compute the new state, then the state gradient must be computed 4 times per time step.

For the MI method, setup only requires forming the FEM matrices and the n_{cp} element state space models. These processes are $O(n)$ and $O(n_{cp})$, respectively. Solving (44) is the major work required to compute an estimate for the state gradient. We consider two alternatives: direct methods using LU factorization; and the conjugate gradient method. Strang (2007) provides an overview of these two methods with general results. We adapt those results to the specifics of the cochlear models considered in this paper.

The work required for LU factorization is dependent on the size n and the *bandwidth* w , or maximum distance of nonzero terms from the main diagonal, of the matrix. For our worst-case estimate, we pessimistically assume bandwidths of $n_x H/L$ in 2D and $n_x^2 HW/L^2$ in 3D. Various algorithms can be used to improve these values but the bandwidth remains

TABLE III. Storage requirements for the largest data structures in the 3 methods. \mathbf{A} is required by both FE MSS and FE MSSV during execution. They also require \mathbf{K} and $\mathbf{K} - \mathbf{R}\mathbf{C}_E\mathbf{B}_E\mathbf{Q}$, respectively, to form \mathbf{A} . FE MI only requires $\mathbf{K} - \mathbf{R}\mathbf{C}_E\mathbf{B}_E\mathbf{Q}$. Worst case figures quoted are the estimated number of double precision floating point values that must be stored for the scenario outlined in the text.

Data Structure	Order	Worst Case	
		2D	3D
\mathbf{K} or $\mathbf{K} - \mathbf{R}\mathbf{C}_E\mathbf{B}_E\mathbf{Q}$	$O(n_{cp}^D)$	8.1×10^5	1.458×10^8
\mathbf{A}	$O(n_{cp}^2)$	1.44×10^8	

$O(n_{cp})$ for 2D and $O(n_{cp}^2)$ for 3d. Factorization takes approximately $2nw^2$ operations which is 6.48×10^8 and 1.4×10^{14} operations in 2D and 3D for our worst-case example [$O(n_{cp}^4)$ and $O(n_{cp}^7)$, respectively]. (The time for these operations will also follow these trends.) LU factorization also requires additional storage space of $O(nw)$, which is $O(n_{cp}^3)$ in 2D and $O(n_{cp}^5)$ in 3D.

Once the LU factorization is complete, it can be used multiple times to solve the system provided that terms in $\mathbf{K} - \mathbf{R}\mathbf{C}_E\mathbf{B}_E\mathbf{Q}$ do not change. Solving the system using the factorization requires approximately $4nw$ operations which is $O(n_{cp}^3)$ in 2D and $O(n_{cp}^5)$ in 3D. For our worst-case example solving the system requires 2.2×10^7 and 7.8×10^{10} operations. As noted in Strang (2007), direct solvers using LU factorization are competitive in 2D even for relatively large problems but in 3D only small problems are feasible. While preparing the results section we used UMFPack (Davis, 2004), a sparse-direct solver, extensively.

For large systems where direct methods are not feasible the conjugate gradient method (CG) becomes the method of choice if the matrix is symmetric and positive definite. Both \mathbf{K} and $\mathbf{K} - \mathbf{R}\mathbf{C}_E\mathbf{B}_E\mathbf{Q}$ satisfy these requirements. CG is an iterative method where a guess is refined successively by estimating the error in matrix equation for that guess. Each iteration requires a multiplication operation on the matrix, which is $O(n_{cp}^D)$. For matrices with condition number κ CG converges to a specified tolerance in less than $O(\sqrt{\kappa})$ iterations. Therefore, the condition numbers discussed in Appendix B are once again important, but they can also be improved through the use of preconditioners to aid convergence. We do not provide estimates for the overall order of CG with respect to n_{cp} because of the large number of factors affecting the computation.

For the MSS and MSSV methods, computing the state gradient requires Eq. (1). Multiplication of $\mathbf{A}\mathbf{x}(t)$, which requires $O(n_{cp}^2)$ operations (2.88×10^8 for the worst-case system), dominates the work so this computation is independent of dimension.

MSSV and MSS both need the inverse of a sparse matrix to form the final state space matrices. The actual inversion should be avoided and either LU factorization or CG⁶ used instead, applied to the columns of \mathbf{R} individually. This process is equivalent to computing $n_{cp}/4$ update steps with the MI method if the RK4 algorithm is used.

The MSS method requires an additional step where $[\mathbf{I} - \mathbf{B}_E\mathbf{Q}\mathbf{K}^{-1}\mathbf{R}\mathbf{C}_E]^{-1}$ is computed. This matrix is full and nonsymmetric so CG cannot be used to invert this matrix and LU factorization requires $2/3n_{cp}^3n_{states}^3$, $O(n_{cp}^3)$, operations. For our worst-case example this is 1.15×10^{12} operations. The zero terms that we noted in \mathbf{A} but commented were not useful at that stage can now reduce this work required. For the system studied \mathbf{C}_E and \mathbf{B}_E are only nonzero for the first state associated with each element. Therefore, $\mathbf{B}_E\mathbf{Q}\mathbf{K}^{-1}\mathbf{R}\mathbf{C}_E$ scatters the $n_{cp} \times n_{cp}$ $\mathbf{Q}\mathbf{K}^{-1}\mathbf{R}$ matrix in an $n_{cp}n_{states} \times n_{cp}n_{states}$ space. By suitable permutations of rows and columns in the system, two nonzero blocks can be collected. The first is $\mathbf{I} - \mathbf{Q}\mathbf{K}^{-1}\mathbf{R}/M$, which requires just $2/3 n_{cp}^3$ or 1.8×10^{10} operations for LU factorization, and an identity block. In this case linear algebra libraries are able to make use of this property.

TABLE IV. Major areas where computational effort is required in the three methods. Factoring matrices dominates the initialization stage of the methods. FE MI and FE MSSV only require $\mathbf{K} - \mathbf{RC}_E \mathbf{B}_E \mathbf{Q}$ to be factored. FE MSS requires both \mathbf{K} and $\mathbf{I} - \mathbf{B}_E \mathbf{Q} \mathbf{K}^{-1} \mathbf{RC}_E$ to be factored. Estimating the state gradient requires $\mathbf{K} - \mathbf{RC}_E \mathbf{B}_E \mathbf{Q}$ to be solved for FE MI and the multiplication \mathbf{Ax} for FE MSS and FE MSSV. Worst case figures quoted are the estimated number of double precision floating point operations that must be performed for the scenario outlined in the text. Only values for LU factorization are presented in this table. The high number of operations required in 3D motivates use of iterative solvers such as the conjugate gradient method.

Data Structure	Order		Worst Case	
	2D	3D	2D	3D
factor \mathbf{K} or $\mathbf{K} - \mathbf{RC}_E \mathbf{B}_E \mathbf{Q}$	$O(n_{cp}^3)$	$O(n_{cp}^7)$	6.48×10^8	1.4×10^{14}
factor $\mathbf{I} - \mathbf{B}_E \mathbf{Q} \mathbf{K}^{-1} \mathbf{RC}_E$	$O(n_{cp}^3)$		1.15×10^{12}	
solve \mathbf{K} or $\mathbf{K} - \mathbf{RC}_E \mathbf{B}_E \mathbf{Q}$	$O(n_{cp}^3)$	$O(n_{cp}^5)$	2.2×10^7	7.8×10^{10}
\mathbf{Ax}	$O(n_{cp}^2)$		2.88×10^8	

This property arises because the first state is associated with the basilar membrane velocity. In the model, the fluid pressure only affects the basilar membrane acceleration directly and this same acceleration is the feedback to the fluid model. Models that have the same input and output pathways will also offer this opportunity to reduce computational complexity.

Table IV provides a summary of the order numbers and figures presented in this section. The computations show that the MI method is more efficient than both the MSS and MSSV for 2D models, both in terms of the time required to initialize the problem and to compute a state update. This assumes that the LU factorization of $\mathbf{K} - \mathbf{RC}_E \mathbf{B}_E \mathbf{Q}$ occurs once. In 3D models the expense of assembling \mathbf{A} explicitly begins to pay for itself and the MSS and MSSV methods should out perform the MI method for large 3D models when updating the state. For our 3D worst-case example, a single step of CG is as expensive as evaluating \mathbf{Ax} because the matrices are a similar size in memory. However, if nonlinear terms in the model cause terms in the \mathbf{A} matrix to change, then the extra work spent on assembling it might not be justified. The MSSV method does not require a second matrix inversion like the MSS method implying that the \mathbf{A} matrix requires less time to construct using MSSV.

¹“Stepwise” could possibly be a more accurate designation but it was not adopted because the resulting acronym (MS) would be easily confused with MSS as introduced below.

²State space has been applied in other cochlear models to represent elements of the cochlear partition [Liu and Neely (2010) is a recent example]. More rarely it has been used to represent the whole model (Diependaal and Viergever, 1983). Overall, state space is a flexible technique that extends beyond its use in this context. We would therefore prefer a designation that specifically reflects their contribution.

³The authors would like to thank Ben Lineton for providing the height value, and confirming that the values of C_3 and C_4 described are also required.

⁴We refer to models following Kim and Xin (2005) in this manner since a number of authors have recently adopted the term, but we note that there is potential for confusion with the term “locally active,” defined in de Boer (1996).

⁵This rule assumes that the matrices have elements of a similar size.

⁶It is uncommon to use CG in this way. Usually CG is applied to $\mathbf{Kp} = \mathbf{b}$, where \mathbf{p} is the unknown and \mathbf{b} is the “right hand side” of the equation. Note that for our problem $\mathbf{b} \in \text{span}(\mathbf{R})$, therefore $\mathbf{b} = \alpha_1 \mathbf{R}_1 + \dots + \alpha_{n_{cp}} \mathbf{R}_{n_{cp}} + \mathbf{1}$. Neglecting errors due to floating point calculations, CG is linear in the sense that the estimated solution and error after m iterations for $\alpha_i \mathbf{R}_i$ is α_i times the value for \mathbf{R}_i . Therefore, if CG can be used to compute \mathbf{p} for $\mathbf{b} = \mathbf{R}_1 + \dots + \mathbf{R}_{n_{cp}} + \mathbf{1}$ using columns of \mathbf{R} separately, then it can be used for any time step of a simulation. However, a different tolerance will probably be required for a given total error in the solution if the columns of \mathbf{R} are used instead of solving \mathbf{b} directly.

Allen, J. (2001). “Nonlinear cochlear signal processing,” in *Physiology of the Ear*, 2nd ed. (Singular Thomson Learning, San Diego, CA), pp. 393–442.

Allen, J. B., and Sondhi, M. M. (1979). “Cochlear macromechanics: Time domain solutions,” *J. Acoust. Soc. Am.* **66**, 123–132.

Bangerth, W., Hartmann, R., and Kanschat, G. (2007). “deal.II—a general-purpose object-oriented finite element library” *ACM Trans. Math. Softw.* **33**, 1–27.

Bangerth, W., Hartmann, R., and Kanschat, G. (2008). “deal.II differential equations analysis library, technical reference,” <http://www.dealii.org> (Last viewed April 8, 2011).

Bertaccini, D., and Sisto, R. (2011). “Fast numerical solution of nonlinear nonlocal cochlear models,” *J. Comp. Phys.* **230**, 2575–2587.

Cai, H., Manoussaki, D., and Chadwick, R. (2005). “Effects of coiling on the micromechanics of the mammalian cochlea,” *J. R. Soc. Interface* **2**, 341–348.

Cheney, W., and Kincaid, D. (2008). *Numerical Mathematics and Computing*, 6th ed. (Thomson Brooks/Cole, Belmont, CA), pp. 1–763.

Davis, T. A. (2004). “Algorithm 832: Umfpack v4.3—An unsymmetric-pattern multifrontal method,” *ACM Trans. Math. Software* **30**, 196–199.

de Boer, E. (1980). “Auditory physics. physical principles in hearing theory. I,” *Phys. Rep.* **62**, 87–174.

de Boer, E. (1991). “Auditory physics. physical principles in hearing theory. III,” *Phys. Rep.* **203**, 125–231.

de Boer, E. (1996). “Mechanics of the cochlea: Modeling efforts,” *The Cochlea*, edited by P. Dallos, A. B. Popper, and R. R. Ray (Springer-Verlag, New York), pp. 258–317.

de Boer, E., and Viergever, M. A. (1982). “Validity of the Liouville-Green (or WKB) method for cochlear mechanics,” *Hearing Res.* **8**, 131–155.

Diependaal, R. J., Duifhuis, H., Hoogstraten, H. W., and Viergever, M. A. (1987). “Numerical methods for solving one-dimensional cochlear models in the time domain,” *J. Acoust. Soc. Am.* **82**, 1655–1666.

Diependaal, R. J., and Viergever, M. A. (1983). “Nonlinear and active modelling of cochlear mechanics: A precarious affair,” in *Mechanics of Hearing: Proceedings of the IUTAM/ICA Symposium* (Martinus Nijhoff, Leiden, Netherlands), pp. 153–160.

Diependaal, R. J., and Viergever, M. A. (1989). “Nonlinear and active two-dimensional cochlear models: time-domain solution,” *J. Acoust. Soc. Am.* **85**, 803–812.

Duifhuis, H., Hoogstraten, H., van Netten, S., Diependaal, R., and Bialek, W. (1986). “Modelling the cochlear partition with coupled Van der Pol Oscillators,” in *Peripheral Auditory Mechanisms* (Springer, Berlin), pp. 290–297.

Elliott, S. J., Ku, E. M., and Lineton, B. (2007). “A state space model for cochlear mechanics,” *J. Acoust. Soc. Am.* **122**, 2759–2771.

Fish, J., and Belytschko, T. (2007). *A First Course in Finite Elements* (Wiley, Chichester, England), pp. 1–319.

Geisler, C. D., and Sang, C. (1995). “A cochlear model using feed-forward outer-hair-cell forces,” *Hearing Res.* **86**, 132–146.

Henderson, H. V., and Searle, S. R. (1981). “On deriving the inverse of a sum of matrices,” *SIAM* **23**, 53–60.

How, J. A., Elliott, S. J., and Lineton, B. (2010). “The influence on predicted harmonic and distortion product generation of the position of the nonlinearity within cochlear micromechanical models,” *J. Acoust. Soc. Am.* **127**, 652–655.

Janssen, M. J. A., Segal, A., and Viergever, M. A. (1978). “Finite element solution of a two-dimensional cochlear model,” *J. Acoust. Soc. Am.* **64**, S133–S133.

Jones, E., Oliphant, T., Peterson, P., et al. (2001). “SciPy: Open source scientific tools for Python,” <http://www.scipy.org> (Last viewed April 8, 2011).

Kagawa, Y., Yamabuchi, T., Watanabe, N., and Mizoguchi, T. (1987). “Finite element cochlear models and their steady state response,” *J. Sound Vib.* **119**, 291–315.

- Kim, D. O., Molnar, C. E., and Matthews, J. W. (1980). "Cochlear mechanics: Nonlinear behavior in two-tone responses as reflected in cochlear-nerve-fiber responses and in ear-canal sound pressure." *J. Acoust. Soc. Am.* **67**, 1704–1721.
- Kim, Y., and Xin, J. (2005). "A two-dimensional nonlinear nonlocal feed-forward cochlear model and time domain computation of multitone interactions." *Multiscale Model. Simul.* **4**, 664–690.
- Kolston, P. J., and Ashmore, J. F. (1996). "Finite element micromechanical modeling of the cochlea in three dimensions." *J. Acoust. Soc. Am.* **99**, 455–467.
- Ku, E. M., Elliott, S. J., and Lineton, B. (2009). "Limit cycle oscillations in a nonlinear state space model of the human cochlea." *J. Acoust. Soc. Am.* **126**, 739–750.
- Lim, K.-M., and Steele, C. R. (2002). "A three-dimensional nonlinear active cochlear model analyzed by the WKB-numeric method." *Hearing Res.* **170**, 190–205.
- Liu, Y.-W., and Neely, S. T. (2010). "Distortion product emissions from a cochlear model with nonlinear mechano-electrical transduction in outer hair cells." *J. Acoust. Soc. Am.* **127**, 2420–2432.
- Matthews, J. W. (1980). "Mechanical modeling of nonlinear phenomena observed in the peripheral auditory system." Ph.D. thesis, Washington University, St. Louis, MO.
- Moleti, A., Paternoster, N., Bertaccini, D., Sisto, R., and Sanjust, F. (2009). "Otoacoustic emissions in time-domain solutions of nonlinear non-local cochlear models." *J. Acoust. Soc. Am.* **126**, 2425–2436.
- Neely, S. T. (1981). "Finite difference solution of a two-dimensional mathematical model of the cochlea." *J. Acoust. Soc. Am.* **69**, 1386–1393.
- Neely, S. T., and Kim, D. O. (1986). "A model for active elements in cochlear biomechanics." *J. Acoust. Soc. Am.* **79**, 1472–1480.
- Parthasarathi, A. A., Grosh, K., and Nuttall, A. L. (2000). "Three-dimensional numerical modeling for global cochlear dynamics." *J. Acoust. Soc. Am.* **107**, 474–485.
- Puria, S., and Allen, J. B. (1991). "A parametric study of cochlear input impedance." *J. Acoust. Soc. Am.* **89**, 287–309.
- Sisto, R., Moleti, A., Paternoster, N., Botti, T., and Bertaccini, D. (2010). "Different models of the active cochlea, and how to implement them in the state-space formalism." *J. Acoust. Soc. Am.* **128**, 1191–1202.
- Steele, C. R., and Taber, L. A. (1979). "Comparison of WKB calculations and experimental results for three-dimensional cochlear models." *J. Acoust. Soc. Am.* **65**, 1007–1018.
- Strang, G. (2007). *Computational Science and Engineering* (Wellesley-Cambridge Press, Wellesly, MA), Chap. 7.
- van Hengel, P. W., Duifhuis, H., and van den Raadt, M. P. (1996). "Spatial periodicity in the cochlea: The result of interaction of spontaneous emissions?" *J. Acoust. Soc. Am.* **99**, 3566–3571.
- van Hengel, P. W. J. (1996). "Emissions from cochlear modelling." Ph.D. thesis, University of Groningen, Groningen, Netherlands.
- Viergever, M. (1980). "Mechanics of the inner ear." Ph.D. thesis, Delft University of Technology, Delft, Netherlands.
- Viergever, M. A. (1977). "A two-dimensional model for the cochlea ii. the heuristic approach and numerical results." *J. Eng. Math.* **11**, 11–28.
- Watts, L. (2000). "The mode-coupling Liouville-Green approximation for a two-dimensional cochlea model." *J. Acoust. Soc. Am.* **108**, 2266–2271.

Appendix D

Investigations into Time Stepping Methods for Cochlear Models

This appendix reproduces the article [Rapson and Tapson \(2011b\)](#) that was submitted to the International Journal of Numerical Methods in Biomedical Engineering. The text presented represents the corrected article that was resubmitted to the journal, and hence uses the template of that journal. Since this thesis was examined, that journal has indicated that they will not publish the paper. (It was originally submitted to the proceedings of AfriCOMP 2011, carried in that journal, but missed the deadline due to the corrections required.) While the content of the paper is still relevant to this thesis, it is superseded by an upcoming manuscript. Therefore, it will only be published in this appendix. My coauthor, Jonathan Tapson, gave his consent for me to include the article verbatim. This paper also represents my own work for the purposes of assessment of this thesis.

Investigations into Time Stepping Methods for Cochlear Models[†]

M. J. Rapson^{1*} and J. C. Tapsen¹

¹*Department of Electrical Engineering, University of Cape Town, Cape Town, South Africa*

SUMMARY

The cochlea, which forms part of the inner ear, is responsible for transducing mechanical sound input into electrical neural impulses. It attracts interest from scientists and engineers in a wide range of disciplines. Cochlear physiology is characterised by strong fluid-structure interaction (FSI) and nonlinear behaviour. Cochlear models have proven their worth as a tool to improve our understanding of the cochlea. The size of various cochlear structures suggests that multiscale models will become important. The majority of cochlear models rely on a specific time stepping method introduced in 1979. Phrased in the current language of FSI, this method is monolithic. In this paper, we examine a potential instability (discussed in the literature) in the standard time stepping scheme applied to cochlear models, and then test a sequentially staggered partitioned scheme. The scheme was designed within *a priori* specified constraints. The investigation into the instability reaffirms that the monolithic system is adequate for the class of cochlear models to which it has been applied. The sequentially staggered partitioning scheme showed some promise, but ultimately produces poor results under the constraints imposed. However, the insight into why the sequentially staggered scheme fails can assist in the development of future multiscale cochlear models. Copyright © 2011 John Wiley & Sons, Ltd.

Received . . .

KEY WORDS: Cochlear Model; Hopf Bifurcation; Fluid-Structure Interaction

1. INTRODUCTION

The cochlea is the part of the ear responsible for converting sound waves into neural signals to be carried to the brain. Our sense of hearing is remarkable. Humans can hear sounds from about 20 Hz to 20 kHz, a range of approximately 10 octaves. We can also hear sounds that vary in intensity by 120 dB sound pressure level. The most sensitive ears can hear sounds below 0 dB or $2 \times 10^{-5} \text{ Nm}^{-2}$. These figures are exceeded by numerous other mammals, some of which have developed extreme sensitivity over a small frequency range – a so-called auditory fovea – e.g. in echolocating bats.

There are a number of reasons to study the cochlea. It is central to our hearing range and performs significant processing of the auditory signal before it reaches the brain. Perhaps the most obvious reason to study the cochlea is to help prevent and improve treatment of hearing loss. History shows that this has been effective. The mechanisms of hearing loss are much better understood now [1], and the research has allowed improvements to hearing aid design as well as hearing tests for babies within hours of birth [2]. A quite separate reason to study the cochlea is to attempt to emulate its capabilities in new technologies. Bio-inspired technologies

*Correspondence to: Rm 3.33, Department of Electrical Engineering, University of Cape Town, Private Bag, Rondebosch, 7701, Cape Town, South Africa. E-mail: michael.rapson@uct.ac.za

[†]Originally presented at AfriCOMP 2011, entitled "Emulating the Theoretical Bifurcation Diagram of a Cochlear Model in Computation"

draw on the principles seen in nature in order to avoid “reinventing the wheel”. For example, a thorough understanding of how the cochlea and the brain work together to understand speech could revolutionise human-machine interfaces.

The cochlea has been the subject of study since the 16th century. Much of the anatomy of the cochlea was known by the start of the 19th century, but this century saw rapid improvements in our understanding of the physiology. The cochlea is known to have *active mechanisms* (amplification) and much of the debate and research is centered around the details of these mechanisms. Cochlear modeling is one of the tools available to researchers in the field. They allow the available information to be placed in a framework to examine our understanding and identify gaps or conflicts in the knowledge. Since the cochlea is fragile, inaccessible and easily damaged, they provide a means to predict the behaviour of parts of the cochlea where direct measurement is difficult or impossible. Furthermore, they allow direct and indirect measurements of the cochlea’s behaviour to be compared. Finally, these models form a natural link between the physical cochlea and technologies inspired by it, for example very large scale integration (VLSI) silicon cochleae [3].

This paper anticipates that the reader has a background in numerical methods but not necessarily cochlear modeling. Therefore, sections 2, 3 and 4 aim to provide an overview of cochlear modeling before moving on to the specific type of model addressed in this paper. This overview is tailored to address certain questions we expect an audience familiar with other numerical methods to have and to provide sufficient background for the remainder of the paper. Therefore, certain important parts of the literature are brushed over. Section 3 in particular provides a number of references for each modeling approach discussed. These are provided for readers who are interested in following up on one of the methods mentioned in passing and are not generally required later in the paper.

The goal of this paper is to put the spotlight on the cochlear models themselves. Time domain models are widely used to investigate the active mechanisms, and the consequential nonlinearities, that are present in the cochlea. Section 5 investigates a potential numerical instability in a widely used time domain cochlear modeling approach that was reported in the literature. The majority of time domain models, including the one for which the instability was described, share a common *monolithic* structure. Section 6 motivates why a *partitioned* structure might be useful for future cochlear models and investigates a possible way of constructing such a structure for a simplified cochlear model. Section 7 studies the effect of extending the proposed partitioned method to realistic cochlear models. Our interest in partitioned cochlear models is due to the fact that they will facilitate future multiscale cochlear models. This work also tests, and is tested in, a cochlear modeling framework that we are developing[†].

2. COCHLEAR ANATOMY AND PHYSIOLOGY

In this section we provide a brief overview of the anatomy and physiology of the cochlea. Besides a general overview, we focus on key facts about the cochlea that are important for the rest of the work presented. Readers who are interested in further information are directed to one of the many excellent books on the topic. References [4] and [5] are two books that the authors can recommend and they were consulted for this section.

2.1. Anatomy

The cochlea forms part of the peripheral auditory system. Figure 1 shows a diagram of its parts. Sound waves in air enter the outer ear and are funneled towards the tympanic membrane (ear

[†]The code for this framework will be released shortly under an open source license. Interested parties can contact the corresponding author for details about how to obtain the code once it is released. The specific programs used in this work are available immediately from the corresponding author.




Figure 1. Diagram showing the location of the outer, middle and inner ear as well as cochlear nerve. Adapted and reproduced from [7, pg. 106]. Reproduced by permission of Dr. Kardou.

drum) that separates the outer ear from the middle ear cavity. The three bones in the middle ear (malleus, incus and stapes) transmit vibrations of the tympanic membrane into the fluid-filled cochlea. This process performs impedance matching between the air and the cochlear fluid so that the sound input is delivered to the fluid efficiently.

The cochlea itself lies in close proximity to the vestibular system, which is responsible for our sense of balance; in fact they are in direct fluid contact. The cochlea is approximately 35 mm long in humans, but is coiled like a snail shell for compactness. At the end nearest the vestibular system (called the *base* or *basal* end of the cochlea) are the oval window and the round window. The oval window is where the last bone in the middle ear (the stapes) connects and it serves as the usual location for input to the cochlea. The round window is flexible to allow the incompressible fluid to move when the stapes deflects. The cochlea narrows slightly from the base to the *apex* or *apical* end of the cochlea. The average cross-sectional area of the cochlea is about 2 mm^2 in humans [6].

The function of the cochlea is to convert the mechanical sound input to an electrical neural signal that the higher stages of the auditory system can process further. Therefore, the cochlea is innervated by the cochlear nerve (part of the 8th cranial nerve).

Figure 2 shows a typical cross-section through the cochlea. The cochlea is embedded in the temporal bone, which is particularly dense. The fluid channel is divided into three *scalae*: the scala vestibuli, scala media and scala tympani. The scala vestibuli and scala tympani ultimately join at the apex of the cochlea, at a place called the *helicotrema*. Only a thin membrane separates the scala media from the scala vestibuli. This membrane allows the two fluids to have different ionic concentrations, but to behave as a single fluid body from a mechanical point of view.

The organ of Corti complex that separates the scala tympani from the remaining *scalae* plays a vital role in the cochlea's behaviour. It contains the basilar membrane, which is stiffened by a series of radial fibers. The organ of Corti proper lies on top of the basilar membrane. Amongst

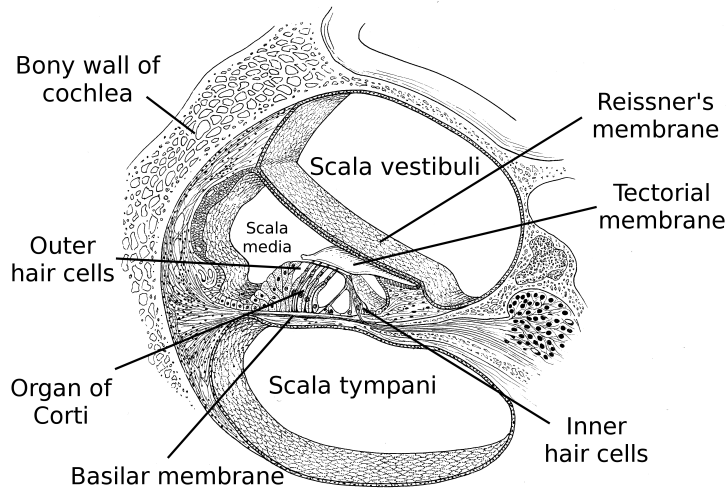


Figure 2. Cross-section through a turn of the cochlea showing the three fluid-filled scalae and the organ of Corti complex. Adapted and reproduced from [7, pg. 109]. Reproduced by permission of Dr. Kardon.

other cells, it contains the inner and outer hair cells, where afferent nerve fibers (towards the brain) arise and efferent fibers (away from the brain) terminate in the cochlea. These cells have characteristic stereocilia that look like hairs and project into the scala media. The stereocilia of the outer hair cells are embedded in the tectorial membrane, which is described as gelatinous, although it also contains a dense network of fibers. Together, all of these parts are sometimes referred to as the *cochlear partition* for convenience.

Various cochlear models assign different functions to each of these structures. Although much is known about their properties and behaviour, there are still areas of debate, particularly centred around the outer hair cells. Given that these cells are embedded in bone, encapsulated by scalae containing fluid and surrounded by the delicate organ of Corti complex, it is understandable that reliable data for the behaviour of the cochlea is hard to obtain. In fact, Georg von Békésy, Nobel Prize winner for his work on the cochlea, has been quoted as saying, “It is probably not a good idea to work on something as complicated and as easy to damage as the organ of Corti.” Modern experimentation techniques have come along way in overcoming these difficulties, but cochlear modeling has an important role to play in integrating the data available.

It is known that the human cochlea contains about 3000 inner hair cells and 3 to 5 outer hair cells for every inner hair cell. There has been some debate over whether the longitudinal stiffness of the cochlear partition is significant. It is generally treated as being negligible, which implies that radial segments containing an inner hair cell and a few outer hair cells are able to move independently of one another, apart from their coupling through the surrounding fluid. The mechanical properties of the organ of Corti complex – mainly the radial stiffness and damping – vary over the length of the cochlea. This causes each of these segments to respond best to its own *characteristic frequency*. In humans these lie between about 20 kHz (at the base) to 20 Hz (at the apex). Alternatively, each segment may be referred to as the *best place* for its characteristic frequency.

2.2. Physiology

In the 1920s [8] Georg von Békésy first observed a *traveling wave* in the cochlea using cadavers and light microscopes. This wave is the result of the interaction of the fluid in the cochlea with

the flexible cochlear partition. It has been the basis of the majority of cochlear models since that time[‡].

By the 1980s it was becoming apparent that the *passive* cochlear models in use were not adequate to describe the behaviour of a healthy cochlea. What were missing were the *active* mechanisms, now considered to reside in the outer hair cells. Active mechanisms use chemical energy to enhance the input sound, and the outer hair cells are a critical part of these mechanisms. The active cochlea shows sharper frequency tuning and lower hearing thresholds. It also shows nonlinear behaviour. Some of the interesting effects that these nonlinear active mechanisms display are:

- Tones interact such that one may *mask* another, or an input containing only two tones will generate a third in the cochlea that is perceived in addition.
- When a signal is played into the ear, a signal is returned from the cochlea that may be measured in the ear canal. This signal is not simply a reflection, but is due to the active processes in the cochlea. This type of phenomenon is referred to as an evoked otoacoustic emission.
- The response of the cochlea shows a compressive nonlinearity. For example, a sound pressure level of 10 dB causes a deflection of the basilar membrane of just 0.35 nm even with the active mechanisms acting to amplify the signal. At 50 dB the basilar membrane deflects about 3.5 nm. Thus the deflection is only 10 times larger while the input signal is 100 times larger. Numbers are for the guinea-pig, from [4, figure 3.10] .

Having introduced the cochlea, and hopefully shown why it is of interest to scientists, we move on to give a brief overview of how the anatomy and physiology are generally translated into cochlear models.

3. COCHLEAR MODELING

The field of cochlear modeling has developed alongside physiological and phenomenological research on the cochlea, mostly as a means for testing the theories on how the cochlea functions. Many cochlear models aim to remain as true to the anatomy and physiology as possible, and this paper considers some of these types of models. Building a model of this type requires that certain key features are well represented. The most important features are the cochlear fluid; cochlear partition; stapes boundary condition and middle ear characteristics; and either the apical or the round window boundary condition.

Linear cochlear models allow the middle ear characteristics to be neglected initially, by treating the stapes as the input to the system and post processing middle ear behaviour. This is not possible with nonlinear models. However, this work deals primarily with a model (proposed in [9]) that was initially solved using a quasilinear technique that required unusual stapes boundary conditions. Since the model that we ultimately investigate neglects the middle ear, it will not be discussed further.

An important distinction in cochlear modeling is between macromechanical and micromechanical models. The former are models of the behaviour of the cochlea as a whole. They consider the fluid structure interaction with multiple radial segments of the cochlear partition, but do not describe the anatomy of the cochlear partition in any detail [10, 11, 12, 13]. The latter make the cochlear partition their subject, for example see references [14, 15]. They attempt to accurately describe the behaviour of the basilar membrane, tectorial membrane and hair cells as separate elements. In practice a cochlear model consists of both, and strikes a balance between detail in the macromechanical model and the micromechanical model, for

[‡]It is important to note that the traveling wave is observed in the motion of the cochlear partition, but arises due to the interaction of the partition and the surrounding fluid. The cochlear partition alone cannot support a wave traveling in the longitudinal direction because the longitudinal coupling within the cochlear partition is negligible.

example [16]. The fact that this tradeoff is already important suggests that multiscale models will naturally find their place in the cochlear modeling landscape, especially since the cochlear length is of the order of 35 mm, a radial segment approximately 10 μm wide, and ionic flows that are critical to the behaviour of hair cells are governed by ion channels that are at most 0.7 nm wide.

Fluid structure interaction is a critical topic in cochlear modeling. It might be surprising then, that the geometry of the fluid is considered fixed, and methods such as mesh moving are not used. As noted above, the displacement of the basilar membrane is of the order of 3.5 nm whereas the cross section of the scalae are about 2 mm², therefore the effects of the change in fluid geometry are generally neglected.

The fluid is generally assumed to be inviscid, irrotational and incompressible. This allows the fluid to be described by the Laplace equation, with either the fluid pressure $p(x, y, t)$ or potential $\phi(x, y, t)$ as the free variable. Neumann boundary conditions on this equation require the gradient of the free variable to be specified. This is proportional to the fluid acceleration in the first case and the fluid velocity (by definition) in the second case. The fluid equation and Neumann boundary conditions are then:

$$\nabla^2 p(x, y, t) = 0 \quad \text{on } \Omega, \quad (1)$$

$$\nabla p(x, y, t) \cdot \vec{n} = -\rho a(x, y, t) \cdot \vec{n} \quad \text{on } \Gamma, \quad (2)$$

or

$$\nabla^2 \phi(x, y, t) = 0 \quad \text{on } \Omega, \quad (3)$$

$$\nabla \phi(x, y, t) \cdot \vec{n} = -v(x, y, t) \cdot \vec{n} \quad \text{on } \Gamma, \quad (4)$$

where Ω is the fluid domain, Γ is the boundary of the fluid domain, ∇ is the gradient operator, ∇^2 is the divergence of the gradient (Laplace) operator, \vec{n} is the outward unit normal, ρ is the fluid density and $a(x, y, t)$ and $v(x, y, t)$ are the acceleration and velocity at a point respectively. We assume a two-dimensional (2D) fluid model in this paper. 1D fluid models are also popular because of their simplicity, although they are known to introduce characteristic artifacts into the solution [17, section 2.3]. 3D models have been used as well [18, 19]. We use the finite element method for spatial discretisation of the fluid model. A wide range of other techniques have been used: finite differences [20]; Green's functions [21]; and recently conformal mappings [22]; as well as field-specific[§] techniques such as transmission line models [25].

The above assumptions have been justified based on the size of the error introduced by neglecting the compressibility, viscosity and nonlinearity of the fluid (see [26] for example). The nonlinearities only become important for simulations of loud sounds (>140 dB). Therefore, this assumption is fair for the full range of interest. Viscosity becomes important for low frequency sound inputs (less than about 100 Hz). Formally, the fluid is viscous but the viscous effects are confined to a thin boundary layer that is neglected. The fact that this approximation breaks down for low frequencies is the source of some numerical difficulties at the helicotrema that need to be treated carefully in time domain models. Puria and Allen [23] found that including viscosity reduced apical reflections, but found that other factors such as scalae tapering can also have a large effect. Compressibility is important for high frequency inputs (formally above about 8 kHz, although it is often neglected even for much higher frequencies). A recent paper found that including compressibility in a frequency domain model [27] improved the agreement between the frequency dispersion in the biological cochlea and in the model after the best place.

[§]A reviewer points out that transmission line models are in some cases equivalent to the 1D finite difference model. This is true, but transmission line models are also used in circumstances that do not correspond neatly to a 1D finite difference model (see [23] and [24] for example). Furthermore, they have strongly influenced how researchers think about the behaviour of cochlear models because they highlight a parallel between the cochlea and electrical transmission line theory. Therefore, we think they are worth acknowledging as a separate, field-specific technique.

The Navier-Stokes equations have been used by some authors [14, 28]. Overall, however, the assumptions have proved their worth in the field.

We will only describe one micromechanical model in detail, as analysed in the rest of the paper. This model was proposed in the 1980s by Stephen Neely [29]. At that time, cochlear models treated the basilar membrane as the dominant element in radial segments and used second order differential equations to describe it as a mass-stiffness-damping system. Neely introduced an additional resonant element that could represent either the tectorial membrane or the stereocilia of the hair cells. This model is still relevant today. [30] used a related model to study otoacoustic emissions.

Whereas Neely included *linear* active mechanisms to represent the outer hair cells, Matthews and Molnar used his model as the starting point for a model with *nonlinear* active mechanisms [31]. We use their model in this work because we wish to investigate the numerical instability that they described. The two-degree-of-freedom system of equations that describes the micromechanics of the cochlear partition is therefore:

$$p_i = M_{1i}a_{1i} + (R_{1i} + R_{3i})v_{1i} + (K_{1i} + K_{3i})d_{1i} - R_{3i}v_{2i} - K_{3i}d_{2i}, \quad (5)$$

$$0 = M_{2i}a_{2i} + (R_{2i} + R_{3i})v_{2i} + (K_{2i} + K_{3i})d_{2i} - R_{3i}v_{1i} - K_{3i}d_{1i}, \quad (6)$$

$$R_{3i}(t) = R_{30}e^{R_{3E}x_i} + (v_{2i}(t) - v_{1i}(t))^2, \quad (7)$$

where M , R , and K represent values for masses, dampings and stiffnesses respectively. a , v , and d represent the acceleration, velocity and displacement respectively. The numerical subscript indicates whether it is a value related to the basilar membrane, 1, hair cell stereocilia, 2, or the coupling between the two, 3. i is an index for the location along the cochlear length because all these parameters and values can vary spatially. Equation (5) represents the dynamics of the basilar membrane, and equation (6) the dynamics of the hair cell stereocilia. Equation (7) is the nonlinear damping term included in their model. R_{30} may be less than zero, in which case it is a negative damping term, which effectively adds energy to the system. Therefore, the system shows (Hopf) bifurcations, as discussed further below.

This style of lumped parameter micromechanical model is popular in the field. It is appropriate because it utilises the available physiological data efficiently. More detailed models have been constructed, for example by setting up a FEM problem for the cochlear partition too [15, 14, 28]. A large number of simplifying assumptions are required where data is not available; and by committing more computational resources to the cochlear partition, these inevitably lead to compromises elsewhere, but this type of model can provide insight into the accuracy of lumped parameter models.

4. CURRENT TIME STEPPING METHODS

To solve a cochlear model in the time domain, the equations discussed in the previous section need to be manipulated into a suitable form. In this section we review the two major time stepping methods that have been applied to this type cochlear model to date. A paper we have recently submitted to the Journal of the Acoustical Society of America (JASA) entitled “Unification and Extension of Monolithic State Space and Iterative Cochlear Models” covers this in further detail, but we present the idea and notation here. In particular, we neglect certain details concerning the stapes input in this section for simplicity because it can be accounted for by the boundary conditions on the fluid domain for the case that we study in the rest of the paper.

The two methods both discretise the time dimension first. This leaves a spatial problem to be solved at each time step. Diependaal *et al.* [32] examined various options for advancing from one time step to the next, and determined that the Runge-Kutta 4th order algorithm (RK4) is optimal for cochlear models of this type. For certain nonlinear models implicit time stepping methods may be required ([33] describes a case where they are).




Figure 3. Diagram showing the rectangular fluid domain Ω with stapes, S , helicotrema, H , and cochlear partition, CP , indicated. Zero normal acceleration is assumed on the upper boundary. The specific boundary conditions assumed for the helicotrema and the stapes will be introduced in section 5.2. The rectangular fluid geometry is symbolic, the finite element method allows other geometries to be used as well. (See section 5.2 for details of the specific geometry used in this paper.) A single FEM cell and two radial segments (the apical segment is highlighted) on the boundary between the fluid and the cochlear partition are marked and shown enlarged on the right. The two radial segments are located at the FEM quadrature points, as suggested by the shape functions indicated. (Linear shape functions were selected for the diagram as the most iconic representation, and do not imply that these must be used.)

4.1. Notation

4.1.1. *FEM Fluid Matrices* FEM was used for the spatial discretisation in this work. The interaction between the cochlear fluid and the cochlear partition is the key feature of a cochlear model. Figure 3 shows the general form of the model. Radial segments are defined to be located at the quadrature points of the cochlear partition boundary Γ_{cp} . This means that the value of the acceleration computed for each segment is the value of the normal acceleration required for the boundary integral. The values of the pressure and normal acceleration in the fluid and cochlear partition are tightly coupled, which means that a change in the acceleration has a large effect on the pressure value and vice versa. The micromechanics of the cochlear model (equations (5)-(7) in this paper) must be satisfied at this boundary.

On the remaining boundaries we need to enforce a value for either the solution (a Dirichlet boundary condition on the pressure) or solution gradient (a Neumann boundary condition). The solution gradient is related to the acceleration by (2). There are various ways to select these values and the specific choices used in this paper are discussed in section 5.2. The two boundary conditions that we will use at the stapes (zero Neumann and zero Dirichlet) both result in an autonomous system.

We convert the strong form of the pressure equations (1-2) to a weak form and replace the continuous spaces by discrete spaces in the usual way (details are available in the aforementioned submission). We depart slightly from the normal approach by manipulating the Gauss quadrature procedure used on the cochlear partition boundary to express the right hand side as a matrix \mathbf{R} multiplied by a vector of normal accelerations at the cochlear partition quadrature points \mathbf{a}_1 , which are also the accelerations of the first degree of freedom of the cochlear partition model. This results in the matrix equation:

$$\mathbf{K}\mathbf{p}_n(t) = \mathbf{R}\mathbf{a}_1(t), \quad (8)$$

where \mathbf{K} is the usual system matrix, $\mathbf{p}_n(t)$ is the pressure at the nodes and \mathbf{R} is a matrix such that $\mathbf{R}\mathbf{a}_1(t)$ is the integral over the cochlear partition that Gauss quadrature would provide. Once again, the subscript 1 indicates that the acceleration of the basilar membrane is required. \mathbf{K} and \mathbf{R} obviously depend on the specific boundary conditions and fluid geometry used. In order to provide feedback of the pressure to the cochlear partition, we define a matrix \mathbf{Q} that

interpolates $\mathbf{p}_n(t)$ to a vector of pressures along the cochlear partition $\mathbf{p}(t)$, located at the quadrature points. We assume that $\mathbf{a}_1(t)$ and $\mathbf{p}(t)$ have the same ordering. Then:

$$\mathbf{p}(t) = \mathbf{Q}\mathbf{p}_n(t). \quad (9)$$

This completes the notation that we will refer to concerning the fluid.

4.1.2. Cochlear Partition State Space Matrices Elliott *et al.* [30] recently introduced a paradigm for studying cochlear models as a single, large state space system. We adopt their paradigm in this paper.

Each radial segment of the cochlear partition (marked in figure 3) can be expressed by a nonlinear state space model with 4 states. A state space model consists of two parts, a set of first order differential equations, the “state update equations” and one or more “output equations”. If the states are defined as $\mathbf{x}_i(t) = [v_{1i} \ d_{1i} \ v_{2i} \ d_{2i}]^T$ at time t the state update equations can be expressed as:

$$\dot{\mathbf{x}}_i(t) = \mathbf{A}_i\mathbf{x}_i(t) + \mathbf{B}_i\mathbf{p}_i(t), \quad (10)$$

$$\mathbf{A}_i = \begin{bmatrix} -\frac{R_{1i}+R_{3i}}{M_{1i}} & -\frac{K_{1i}+K_{3i}}{M_{1i}} & \frac{R_{3i}}{M_{1i}} & \frac{K_{3i}}{M_{1i}} \\ 1 & 0 & 0 & 0 \\ \frac{R_{3i}}{M_{2i}} & \frac{K_{3i}}{M_{2i}} & -\frac{R_{2i}+R_{3i}}{M_{2i}} & -\frac{K_{2i}+K_{3i}}{M_{2i}} \\ 0 & 0 & 1 & 0 \end{bmatrix}, \quad (11)$$

$$\mathbf{B}_i = [1/M_{1i} \ 0 \ 0 \ 0]^T, \quad (12)$$

where $\dot{\mathbf{x}}_i$ is the derivative of the states with respect to time. The $\dot{\cdot}$ notation will be used in this manner. Care must be taken with the R_{3i} term since it is technically nonlinear, which makes a matrix representation of the state space system inappropriate. We will neglect its nonlinear parts in this discussion and reintroduce them at a later stage. In this representation, the first two states relate to the first degree of freedom and the final two states relate to the second degree of freedom.

[30] used a slightly uncommon definition for the output of their state space model: They selected the segment’s velocity as the output rather than the acceleration required by the fluid model. Since this is only a semantic, we follow their definition for easy comparison to earlier papers. Therefore:

$$v_{1i}(t) = \mathbf{C}_i\mathbf{x}_i(t), \quad (13)$$

$$\mathbf{C}_i = [1 \ 0 \ 0 \ 0], \quad (14)$$

$$a_{1i}(t) = \dot{v}_{1i} = \mathbf{C}_i\dot{\mathbf{x}}_i(t). \quad (15)$$

Finally, we collect the individual state space models into a combined state space model, using the same ordering selected for the fluid. The combined matrices \mathbf{A}_E , \mathbf{B}_E , and \mathbf{C}_E all have a block diagonal structure, for example:

$$\mathbf{A}_E = \begin{bmatrix} \mathbf{A}_1 & 0 & \cdots & \\ 0 & \mathbf{A}_2 & & \vdots \\ \vdots & & \ddots & 0 \\ \cdots & 0 & & \mathbf{A}_N \end{bmatrix},$$

see also [30].

4.2. Time Stepping Methods

The first time domain cochlear model was solved by Allen and Sondhi in 1979 [13]. Their approach required that the fluid model be *augmented* with the mass term in their cochlear

partition model. The majority of cochlear modeling papers have followed this general approach. More recently [30] followed a different approach that is only valid if \mathbf{K} is invertible. In the JASA submission mentioned above, we have shown that the two methods are mathematically equivalent. We will follow the method of [13] in this paper since \mathbf{K} may not be invertible, but will use the wholly-state-space representation introduced in [30] for clarity.

Using the collected form of equations (15) and (10) as well as equation (9) we obtain:

$$\mathbf{a}_1(t) = \mathbf{C}_E \mathbf{A}_E \mathbf{x}(t) + \mathbf{C}_E \mathbf{B}_E \mathbf{Q} \mathbf{p}_n(t). \quad (16)$$

This is substituted into equation (8) and terms containing the unknowns \mathbf{p}_n are collected:

$$[\mathbf{K} - \mathbf{R} \mathbf{C}_E \mathbf{B}_E \mathbf{Q}] \mathbf{p}_n(t) = \mathbf{R} \mathbf{C}_E \mathbf{A}_E \mathbf{x}(t). \quad (17)$$

Using equations (10) and (9) again, and collecting $\mathbf{A}_E \mathbf{x}(t)$ terms we obtain:

$$\dot{\mathbf{x}}(t) = \left[\mathbf{I} + \mathbf{B}_E \mathbf{Q} [\mathbf{K} - \mathbf{R} \mathbf{C}_E \mathbf{B}_E \mathbf{Q}]^{-1} \mathbf{R} \mathbf{C}_E \right] \mathbf{A}_E \mathbf{x}(t), \quad (18)$$

which is the state update equation for the combined, fluid coupled system. The output equation is $\mathbf{v}_1(t) = \mathbf{C}_E \mathbf{x}(t)$ if velocities are required. In this work RK4 was used with (18) to advance the simulation time. Rk4 requires four evaluations of $\dot{\mathbf{x}}$ per time step.

There are two approaches to converting (18) to an algorithm. The first is to form the matrix $\left[\mathbf{I} + \mathbf{B}_E \mathbf{Q} [\mathbf{K} - \mathbf{R} \mathbf{C}_E \mathbf{B}_E \mathbf{Q}]^{-1} \mathbf{R} \mathbf{C}_E \right] \mathbf{A}_E$ and use it to compute $\dot{\mathbf{x}}$ for RK4 in the usual manner. The second avoids forming the full matrix by dividing each evaluation of $\dot{\mathbf{x}}$ into a number of substeps: a) compute $\mathbf{R} \mathbf{C}_E \mathbf{A}_E \mathbf{x}$; b) solve (17) for \mathbf{p}_n ; c) use \mathbf{x} and \mathbf{p}_n together with (10) to evaluate $\dot{\mathbf{x}}$. Steps (a) and (c) can be done for each radial segment separately, without forming the combined matrices. These two methods are mathematically equivalent so in the remainder of the paper we use the first approach to analyse our system but will discuss how it may be implemented in the second approach.

We previously neglected the nonlinear terms in the model to allow matrix manipulations to be used freely. These now must be reintroduced into equation (18). The nonlinearities in this specific model are of the form:

$$a_{1i \text{ NL}} = (v_{2i}(t) - v_{1i}(t))^3 / M_{1i}, \quad (19)$$

$$a_{2i \text{ NL}} = -(v_{2i}(t) - v_{1i}(t))^3 / M_{2i}, \quad (20)$$

where the NL subscript shows that we are considering the nonlinear part only. Note that v_{1i} and v_{2i} are states of the system and in particular, they are not dependent on the pressure. This means that the $\left[\mathbf{I} + \mathbf{B}_E \mathbf{Q} [\mathbf{K} - \mathbf{R} \mathbf{C}_E \mathbf{B}_E \mathbf{Q}]^{-1} \mathbf{R} \mathbf{C}_E \right]$ term is unaffected by reintroducing the nonlinearities. The term $\mathbf{A}_E \mathbf{x}(t)$ is affected and this can be handled in two ways: this term can be linearised about v_{1i} and v_{2i} each time an estimate of $\dot{\mathbf{x}}(t)$ is required, and since they are states this is exact; or the nonlinear contribution can be calculated as shown in equations (19) and (20) and added to the linear contribution from $\mathbf{A}_E \mathbf{x}(t)$ before multiplying by the bracketed term. If the nonlinearity affects the bracketed term (\mathbf{B}_E would usually be the site) then more effort is required to incorporate them. (The implicit method proposed in [33] addresses such a case.)

The remainder of the paper divides into three separate but related experiments: section 5, follows up a reference in the literature to a possible numerical problem with the time stepping method just described, that reference described the instability for a simplified cochlear model with just two radial segments; section 6, we investigate a sequentially-staggered partitioning scheme for the 2 segment model; section 7, we test whether the sequentially-staggered partitioning scheme extends to cochlear models with realistic numbers of segments. Sections 5 and 6 were presented at AfriCOMP 2011. Section 7 represents work completed since that conference.




Figure 4. The bifurcation diagram produced by Matthews and Molnar [31, pg. 440] for the cochlear model with 2 segments. The solid line indicates stable stationary solutions; dashed line unstable stationary solutions; filled squares Hopf bifurcation points; and open circles branches of unstable periodic solutions. Reproduced by permission of Plenum Publishers (Springer).

5. A POTENTIAL NUMERICAL PROBLEM?

5.1. Background

We have noted that the approach laid out above has been a popular method for setting up and solving cochlear models since it was introduced in [13]. On one hand this shows that a large number of groups have found it problem free to use, but on the other hand any description of a numerical instability is of concern. Matthews and Molnar used it for a nonlinear cochlear model, and observed instability at values of R_{30} that they expected to be stable. They investigated this further for simplified 2 and 3 segment cochlear models using a nonlinear dynamical systems (NLDS) analysis package and plotted the bifurcation diagram of the system. Figure 4 shows the diagram that they obtained for the 2 segment system. It shows that the bifurcation should occur at $R_{30} = -5.00$ (to two decimal places, for the 2 segment case), this corresponds to $R_{3i} = -R_{2i}$ for the parameters they used. The significance of this is that because $M_1 \gg M_2$ for these parameters, $R_{3i} \approx -R_{2i}$ is required for the second degree of freedom in a segment to affect the first degree of freedom (this is discussed in [29, pg. 104]). They were surprised to find that in their numerical results for a large system with 240 segments the simulations were unstable even for $R_{30} = -4.96$, a value that they thought would make the system stable given the difference between R_{3i} and R_{2i} .

We were particularly interested by this result because it describes an instability around the critical point for a system with a Hopf-bifurcation. In work with collaborators on VLSI silicon cochleae we require the ability to simulate a cochlear model on either side of, and at, a Hopf-bifurcation[¶] accurately (this work is described in [35]). The micromechanical model used in that work is slightly simpler in that it requires only a two state cochlear model but, to understand the instability that was described, we investigate the four state system used by Matthews and Molnar.

[¶]Hendrikus Duifhuis recently showed the relationship between that model and “Van der Pol-oscillator” cochlear models (see [34]).

5.2. Method

The NLDS software used by Matthews and Molnar allowed them to plot the whole bifurcation diagram including unstable solutions. We have been developing a framework for cochlear modeling using a FEM library, deal.II [36, 37] and wanted to use this example as a test of the framework's capabilities. Therefore, we setup equivalent simulations of the 2 segment model in the deal.II cochlear modeling framework and Python using SciPy [38]. These two simulations implement (18). The SciPy simulation used the first approach described in section 4.2 and the deal.II model used the second approach. The deal.II model could be used to simulate the 240 segment model as well as the 2 segment model. Since locating unstable solutions is difficult without specially designed software, we focused on the Hopf bifurcation point as the characteristic feature to emulate and test.

We also determine the stability of the models for specific R_{30} values by examining the eigenvalues of (18). (This is one of the major benefits of the paradigm introduced in [30]. We have recently added this functionality to our deal.II framework.) The SciPy simulation and deal.II framework were benchmarked against one another to verify that they produce identical data.

All simulations used the single scala geometry from [31]. This consists of a rectangular 2D fluid domain 22.5 mm long and 3 mm high. The same parameters as [31] were used. Matthews and Molnar in turn drew the majority of their parameters from [29], parameter set B (PSB). However, they made the masses independent of position on the cochlear partition: $M_1 = 5.0 \times 10^{-3} \text{ gcm}^{-2}$; and $M_2 = 5.0 \times 10^{-7} \text{ gcm}^{-2}$.

Figure 3 shows the location of boundary conditions for the fluid domain. The boundary condition at the helicotrema was zero Neumann (zero flow). Two conditions were tested for the stapes boundary, zero Neumann (zero flow) and zero Dirichlet (zero pressure), described further below. If two Neumann boundary conditions are used, \mathbf{K} is not invertible but equation (18) avoids inverting \mathbf{K} directly.

Both of the possible stapes boundary conditions result in an autonomous system with the trivial solution $\mathbf{x} = \mathbf{0}$, $\mathbf{p}_n = \mathbf{0}$, $t > 0$. Therefore, to start time domain simulations the states of the system are perturbed slightly from zero. From the response of the system to this perturbation we observed whether the R_{30} value was stable (an energy norm that we describe below tended towards zero) or unstable (the norm becomes unbounded). This simple test has an advantage over linear techniques such as "von Neumann stability analysis^{||}" in that it can be used on the nonlinear model directly. We also checked for stability of the origin ($\mathbf{x} = \mathbf{0}$, $\mathbf{p}_n = \mathbf{0}$) by calculating the eigenvalues of the matrix in equation (18). For this calculation, the nonlinear terms can be neglected since they correspond to cubic powers of small numbers at the origin.

5.3. Results

Using the eigenvalues of the system, we found that the bifurcation point of the 2 segment system differed depending on whether the stapes boundary condition was zero pressure (zero Dirichlet) or zero flow (zero Neumann). When zero flow was imposed at that boundary, the bifurcation point was -4.9999994107 (to 10 decimal places), which provides a good match to the bifurcation diagram, figure 4. For zero pressure it was -4.9851285387. Therefore, when the pressure is imposed bifurcation occurs about 0.0149 units earlier. The significance of this result is that it implies the NLDS code imposed the zero flow condition, whereas investigations into the algorithm used for the Green's function model in [31] showed that it imposed zero pressure at the stapes. This suggests that the difference in boundary condition might account for at least part of the instability observed but it does not fully account for instability when $R_{30} = -4.96$.

^{||}This is also, more descriptively, called Fourier stability analysis [39].

We continued by studying the eigenvalues of the 240 segment system using the deal.II model. The bifurcation points were found to be -4.9634** and -4.9634242712 for the zero Neumann and zero Dirichlet boundary conditions respectively. The time domain results from the FEM implementations made for this work are stable for $R_{30} = -4.96$, for both the small and large system and for both the boundary conditions considered. Slight differences between the FEM and Green's functions might cause differences in the bifurcation point of the model. It is possible that the instability they observed was correct behaviour for their implementation because the value of R_{30} that they consider is close to the bifurcation point of the FEM model. The results for the larger system do not show any significant effect of the stapes boundary conditions on the bifurcation point. These results highlight that the condition $R_{3i} \approx -R_{2i}$ only gives the general vicinity of the bifurcation point. For the 2 segment, fixed stapes model it happens to be particularly accurate.

Studying the eigenvalues of the system has verified that the models reproduce the published results (figure 4) and located the bifurcation point for various conditions. Time domain results show that the RK4 provides correct behaviour before, in the vicinity of, and beyond the bifurcation point. Figure 5 shows the time domain response of the 2 segment model for three values of R_{30} : -4.96; -4.99; and -5.01. Results for the 240 segment system are presented at a later stage (see figures 9 and 12). The value of the energy norm E below is plotted on the y-axis:

$$E = \frac{1}{2} \sum_i (M_{1i}v_{1i}^2 + M_{2i}v_{2i}^2 + K_{1i}d_{1i}^2 + K_{2i}d_{2i}^2), \quad (21)$$

where terms containing K_{3i} are neglected because $K_{3i} = 0$ in the parameters considered. This norm neglects the energy contribution of the fluid mass and velocity to the total system energy.

We think that the periodic oscillations visible in the first two plots are due to energy transfer between the segments and the fluid. The period of these oscillations is approximately 375 Hz so they seem to be linked to "whole cochlear resonances": where boundary conditions at the helicotrema cause the reflection of low frequency energy back into the cochlea resulting in a non-physiological sensitivity to these frequencies (these were first described in [40, pg 56] and in [9, pg 70-75]).

The time domain plots for $R_{30} = -4.96$ and -4.99 confirm that these are both stable, as the exponential decay in the amplitude of the energy norm shows. For $R_{30} = -5.01$ the model is unstable and the system reaches a limit cycle. The figure appears to show that the system settles to a first limit cycle at about 10 ms before transitioning to a second limit cycle at 40 ms. The behaviour between 10 ms and 40 ms is governed by equation (6) of the basal segment, and after 40 ms by equation (6) of the apical segment. These modes are individually unstable, which is why the energy norm does not approach zero periodically as it does in the stable case: the fluid is not transferring an appreciable quantity of energy in this case.

The eigenvalue results that we present suggest that the instability described in [31] is not a major concern for the widely used solution method discussed in this paper. Time domain results also support this conclusion that RK4 is adequate for simulating this model.

6. A PARTITIONED METHOD

6.1. Motivation

In the section above we introduced the most common approach to handling the fluid structure interaction and time stepping in the cochlea. We saw that it is entirely adequate for the

**Section 6.4 discusses how the parameter set used has a set of stationary points rather than a single one for the zero Neumann boundary condition. This also showed up in the eigenvalues of the system and disturbed the algorithm used to track the bifurcation point automatically. This value was found by observing the trends in a plot of the number of unstable eigenvalues in the system vs R_{30} hence it is quoted to 4 decimal places only.




Figure 5. The evolution of the cochlear partition energy (see equation (21)) for three values of R_{30} : -4.96 (a); -4.99 (b); -5.01 (c) in a 2 segment model with fixed stapes (zero Neumann).

nonlinear model that Matthews and Molnar addressed, but for some classes of nonlinearities implicit time stepping methods may be required.

The method outlined is *monolithic* in the sense that model is assembled into a single matrix equation that is then solved. *Monolithic* and its opposite *partitioned* are terms used by the fluid-structure interaction community that seem particularly apt to describe cochlear models as well [41]. Furthermore, by using language common in that field we hope to facilitate an exchange of ideas between the two fields. *Partitioned* methods have become popular for fluid-structure interaction problems. They have certain advantages. Notably, they allow smaller and better conditioned matrices to be solved for each of the relevant domains rather than a single large matrix, and they allow state of the art codes and solvers to be used for each domain [42]. Inevitably they have drawbacks as well. In particular they tend to be less stable than a monolithic method and are sometimes more computationally expensive.

A sequentially staggered partitioned system is one where each domain is evaluated only once per time step, alternately. This approach is successful on lightly coupled systems. However, strongly coupled systems require an iteratively staggered partitioned scheme, where the solution approaches the solution to the monolithic system as the number of iterations increases.

The cochlea is known to show strongly coupled fluid-structure interaction behaviour, but we investigate a method to avoid these strongly coupled effects with the aim of designing a sequentially staggered partitioned scheme. We consider a sequentially staggered scheme rather than an iteratively staggered scheme in this paper for two reasons: sequentially staggered schemes are simpler and computationally cheaper when they are suitable so we want to investigate the simpler option before resorting to a more complex solution; and a sequentially staggered scheme is more consistent with the cochlear modeling framework that we are developing because the framework aims to present the user a simplified view of the fluid dynamics and fluid structure interaction. Therefore, we continued the investigations into the system described above and shifted our focus towards a possible alternative to the monolithic solution method.

6.2. Derivation of a partitioned method

With this aim in mind, we first consider the features of equation (18) that make it a successful implementation of this model. The term $\mathbf{K} - \mathbf{R}\mathbf{C}_E\mathbf{B}_E\mathbf{Q}$ consists of the fluid Laplacian matrix *augmented* by the mass terms of the cochlear partition – the same augmentation that was required in the first time domain cochlear model [13]. This is essentially half of a Robin boundary condition as the fluid equation (17) shows.

For the typical mass-spring-damper system:

$$a(t) = \frac{P(t)}{M} - \frac{R}{M} \int a(t) dt - \frac{K}{M} \iint a(t) dt dt. \quad (22)$$

Notice that a step change in pressure causes an instantaneous step response in the acceleration. Similarly, equation (8) shows that for a Laplacian fluid model a step change in the acceleration on the boundary causes an instantaneous step change in the pressure because there are no time dependent terms. Therefore, pressure and acceleration affect each other on a rapid time scale. On the other hand, a step change in the pressure (or acceleration) does not cause a step change in the velocity or displacement. The acceleration is coupled to itself via the velocity and displacement, but this loop is comparatively slow. Therefore, using a Robin boundary condition has effectively separated the fluid simulation into rapid and slow modes of behaviour. All information available about the rapid response of the fluid-structure interface is provided when equation (17) is solved, but the dynamics of the slower response of the cochlear partition are neglected. (Note that the fluid itself has no slower dynamics in this model.) In the time stepping methods discussed above, the slower dynamics are captured by forming the monolithic system equation (18) either implicitly or explicitly. If the simulation of the fluid and the cochlear partition are to be partially decoupled, the rapid dynamics of their interaction must also be captured in the cochlear partition simulation.

The difficulty in handling the fluid is not that slow modes need to be removed, as was the case for the cochlear partition, but that it couples every segment to every other segment. If we captured the full extent of this coupling we would essentially have another monolithic method. Instead, we examine the coupling matrix to determine whether the coupling between some segment pairs can be neglected. From equation (17) we see that:

$$\mathbf{P}(t) = \mathbf{Q} [\mathbf{K} - \mathbf{R}\mathbf{C}_E\mathbf{B}_E\mathbf{Q}]^{-1} \mathbf{R} \mathbf{a}_s(t), \quad (23)$$

$$\mathbf{a}_s(t) = \mathbf{C}_E\mathbf{A}_E\mathbf{x}(t), \quad (24)$$

where $\mathbf{P}(t)$ is the vector of segment pressures and $\mathbf{a}_s(t)$ are the slow dynamics of the acceleration. We examined the 2 segment model with a Neumann boundary condition at the stapes and found:

$$\mathbf{Q} [\mathbf{K} - \mathbf{R}\mathbf{C}_E\mathbf{B}_E\mathbf{Q}]^{-1} \mathbf{R} = \begin{bmatrix} 4.996 \times 10^{-3} & 4.216 \times 10^{-6} \\ 4.216 \times 10^{-6} & 4.996 \times 10^{-3} \end{bmatrix}. \quad (25)$$

The diagonal terms of this matrix are 3 orders of magnitude larger than the off diagonal terms. (This pattern does not hold for the case of a Dirichlet boundary condition but this will be discussed in section 7.) Therefore, the pressure due to the cochlear partition acceleration that a specific segment experiences is most affected by the segments own motion. Figure 6 shows the effect of neglecting the off diagonal terms entirely when simulating the cochlear partition. There is close similarity between the two figures, but errors are apparent after some time and at low values of total energy.

In practice, we do not want to remove the coupling between the segments entirely but instead to partition the response into a rapid component that includes only self coupling and a slower fully coupled response. We therefore define:

$$\mathbf{P}(t) = \mathbf{P}_f(t) + \mathbf{P}_s(t), \quad (26)$$

where $\mathbf{P}_f(t)$ and $\mathbf{P}_s(t)$ are the fast and slow components respectively. We approximate the fast component by the segment self coupling:

$$\mathbf{P}_f(t) \approx \mathbf{P}_{sc}(t) = \mathbf{G} \mathbf{a}_s(t), \quad (27)$$

$$\mathbf{G}_{jk} = (\mathbf{Q} [\mathbf{K} - \mathbf{R}\mathbf{C}_E\mathbf{B}_E\mathbf{Q}]^{-1} \mathbf{R})_{jk} \delta_{jk} \quad (28)$$

where $\mathbf{P}_{sc}(t)$ is the pressure component due to the segment self coupling and δ_{jk} is the Kronecker delta, no summation over ij . Note that $\mathbf{P}_{sc_i}(t)$ depends only on $\mathbf{a}_{s_i}(t)$ and not the acceleration of any other segment. We substitute this approximation, (27), into (26) to obtain the approximation of the slow components:

$$\mathbf{P}_s(t) = \mathbf{P}(t) - \mathbf{P}_{sc}(t). \quad (29)$$




Figure 6. A comparison of the monolithic solution to the 2 segment system with Neumann boundary conditions (a) and the same system with off diagonal terms of $\mathbf{Q}[\mathbf{K} - \mathbf{R}\mathbf{C}_E\mathbf{B}_E\mathbf{Q}]^{-1}\mathbf{R}$ neglected (b). Once again, the partition energy is calculated using equation (21). The parameters for the model are identical to those used in figure 5.

We therefore have a definition for a fast pressure component, dependent only on the local states, and a slow pressure component that captures the behaviour of the remainder of the model. Equation (26) allows us to recover the original pressure and up to this point we have not changed the system in any way.

$\mathbf{P}(t)$ is the solution to the augmented fluid problem. Therefore, every time that $\mathbf{P}_s(t)$ is computed the fluid problem must be formed and solved. (This is relatively cheap for a Laplacian fluid model.) We restrict this operation to once per time step and (26) will use the most recent estimate for the slow component of the pressure, $\tilde{\mathbf{P}}_s(t)$. Then (26) becomes:

$$\mathbf{P}(t) = \mathbf{P}_{sc}(t) + \tilde{\mathbf{P}}_s(t), \quad (30)$$

which is identical to (26) once per time step, but contains an error when used for other estimates.

Using (30) and (10) gives the equation required to update the cochlear partition states:

$$\dot{\mathbf{x}}_i(t) = \mathbf{A}_i\mathbf{x}_i(t) + \mathbf{B}_i\mathbf{P}_{sc_i}(t) + \mathbf{B}_i\tilde{\mathbf{P}}_{s_i}(t). \quad (31)$$

Substituting for \mathbf{P}_{sc} and collecting terms yields:

$$\dot{\mathbf{x}}_i(t) = (\mathbf{I} + \mathbf{B}_i\mathbf{G}_{ii}\mathbf{C}_i)\mathbf{A}_i\mathbf{x}_i(t) + \mathbf{B}_i\tilde{\mathbf{P}}_{s_i}(t). \quad (32)$$

In this case, the term $(\mathbf{I} + \mathbf{B}_i\mathbf{G}_{ii}\mathbf{C}_i)$ is related to the augmentation of the partition with the fluid coupling. The double subscripted \mathbf{G}_{ii} reduces to a single number, where as \mathbf{B}_i and \mathbf{C}_i are the state space matrices for the i^{th} segment.

6.3. Applying the method

The previous section has discussed how the interaction of the fluid and structure can be separated into fast and slow components or local and nonlocal components to obtain an approximation of the state gradient $\dot{\mathbf{x}}_i(t)$. We want to apply the RK4 to the system using this approximation. Recall that it requires four estimates of $\dot{\mathbf{x}}$ and the current states \mathbf{x}_n to estimate the next states \mathbf{x}_{n+1} . Equation (32) provides an approximation of $\dot{\mathbf{x}}$, given an approximation of the slow component $\tilde{\mathbf{P}}_s$ and the states \mathbf{x} .

The steps required to advance from \mathbf{x}_n to \mathbf{x}_{n+1} method are:

1. Using \mathbf{x}_n , solve the fluid equation (17) to obtain the current pressure \mathbf{P}_n . This is the only time (17) needs to be solved during this time step.




Figure 7. The evolution of the cochlear partition energy (see equation (21)) in the partitioned and monolithic model for three values of R_{30} : -4.96 (a); -4.99 (b); -5.01 (c). The model parameters are as for figure 5, in particular this is a 2 segment model.

2. Using \mathbf{x}_n , \mathbf{P}_n and equations (24) and (27), compute \mathbf{P}_{sc} . (This process only requires the local states for each segment.)
3. Using \mathbf{P}_n , \mathbf{P}_{sc} and (29), compute the slow pressure component \mathbf{P}_s . This value of \mathbf{P}_s is used for the entire time step.
4. Apply the RK4 to (32) to integrate \mathbf{x}_n to \mathbf{x}_{n+1} using $\tilde{\mathbf{P}}_s = \mathbf{P}_s$ from (3) and the estimates of \mathbf{x} from the RK4 algorithm itself.

This series of steps can be implemented without forming global matrices \mathbf{A}_E , \mathbf{B}_E and \mathbf{C}_E . An equivalent using global matrices can be implemented (one was used to for testing purposes in this work). The matrix equation for the slow pressure becomes:

$$\mathbf{P}_s = (\mathbf{Q}[\mathbf{K} - \mathbf{R}\mathbf{C}_E\mathbf{B}_E\mathbf{Q}]^{-1}\mathbf{R})_{jk}(1 - \delta_{jk})\mathbf{C}_E\mathbf{A}_E\mathbf{x}(t), \quad (33)$$

with no summation over ij . In practice there would be little point in using the partitioned method with a global matrices because if the inversion required can be performed then the matrix operations required to implement the monolithic approach with global matrices will also be feasible and that approach would be preferred.

We have already noted that because \mathbf{P}_s is computed from \mathbf{x}_n , it is exact for the first RK4 step because that step also uses \mathbf{x}_n . (This obviously assumes that \mathbf{x}_n is accurate.) For the remainder of the RK4 steps \mathbf{P}_s (and hence the total pressure, $\mathbf{P}_s + \mathbf{P}_{sc}$) contains an error.

6.4. Results

The partitioned method was tested in the time domain in the same manner as the monolithic system. (It does not lend itself to eigenvalue analysis as neatly the monolithic system since in this case RK4 and time step size affect stability.) The same fluid geometry and partition parameter were used. In this section only results from a 2 segment model are discussed.

Figure 7 shows the results that were obtained for $R_{30} = -4.96, -4.99$ and -5.01 . The agreement between the partitioned method and the monolithic system shows an improvement over the decoupled system. However, the agreement still becomes worse towards the end of the time period considered. Surprisingly, the proximity of the system to the bifurcation does not seem to be a major factor that affects the performance. The unstable case actually shows the closest correspondence between the two systems, but this is probably due to the fact that in this case the fluid coupling effects are negligible.

The partitioned system is more sensitive to the time step size. For a suitably small time step, the errors introduced by approximating \mathbf{P}_s are negligible and the partitioned system approaches the monolithic system. In this sense it is “stiffer” than the monolithic system. This point will be picked up in section 7 but we first give our attention to whether the data set used for this problem is a fair test of the system.

The data used by Matthews and Molnar was essentially taken from [29], parameter set B (PSB), with the segment mass set to be a constant spatial function. In the PSB, $K_{3i} = 0$ and this removes the coupling due to d_{1i} and d_{2i} between the degrees of freedom represented by equations (5) and (6). If, in addition, a Neumann boundary condition is imposed at the stapes, the system does not have a single steady state solution. In fact it admits solutions of the form:

$$p_i = C, \quad (34)$$

$$v_{1i} = v_{2i} = d_{2i} = 0, \quad (35)$$

$$d_{1i} = C/K_{1i}. \quad (36)$$

Where C is an arbitrary constant. This is because the solution to Laplace equation is only determined up to a multiplicative constant unless a Dirichlet boundary condition is imposed, so any constant pressure is a solution to the fluid equation. In fact, this solution structure shows itself in the time domain results. All of the systems with Neumann boundary conditions that we have presented so far have a very slow unstable mode, the effects of which can just be seen around the 140 ms mark. This unstable mode does not show up in the eigenvalues of the 2 segment system, but was visible in the eigenvalues of the 240 segment system: we put this down to the fact that the system matrix for the 2 segment model is better conditioned than the 240 segment model. We interpret this slow unstable mode as an effect of the multiple stable steady state solutions.

If we impose a Dirichlet boundary condition at the stapes then we restrict the steady state solution to $p_i = 0$ and $\mathbf{x} = \mathbf{0}$, but this destroys the assumed diagonal structure of $\mathbf{Q}[\mathbf{K} - \mathbf{R}\mathbf{C}_E\mathbf{B}_E\mathbf{Q}]^{-1}\mathbf{R}$. Therefore, the results up to this point suggest that the partitioned method shows promise, but it must be tested on a system with more segment to verify whether the 2 segment system is a representative case.

7. PARTITIONED METHOD AND LARGE COCHLEAR MODELS

In section 6 above we presented an approach to partitioning the computation of a cochlear model. The results showed mixed success for a very small system. This might be a negative result since cochlear models regularly contain hundreds to thousands of segments; however, the system allowed a set of solutions rather than just a single solution so it is perhaps not surprising that the partitioned method converges to a slightly different solution than the monolithic. In this section we put the partitioned method to the test on large systems, and discuss its possible extension to other fluid models.

7.1. Structure of the fluid coupling matrix

The $\mathbf{Q}[\mathbf{K} - \mathbf{R}\mathbf{C}_E\mathbf{B}_E\mathbf{Q}]^{-1}\mathbf{R}$ matrix of models with large numbers of segments follow the general pattern observed in the 2 segment case: that values on the main diagonal are larger than those off the diagonal. Figure 8 shows two plots of the matrix, in which each element is represented by a coloured block, on a grey scale. The main diagonal is the prominent feature visible in the plot of the whole matrix on the left hand side. However, focusing on the first ten elements (right hand side) shows that elements on the nearby minor diagonals are of similar orders of magnitude in this case, suggesting that neglecting them may cause significant errors. Nevertheless we proceeded to implement and evaluate the performance of the partitioned system. Interestingly, the Dirichlet boundary condition applied at the 0th segment has a visible effect on itself and the 1st segment but relatively little effect on the remaining segments.

7.2. The effect of Dirichlet boundary conditions

We first return to the question of whether the small system discussed above is a representative or unrepresentative test of the system. As noted above, when Neumann boundary conditions




Figure 8. Figure showing the relative sizes of terms in the $\mathbf{Q}[\mathbf{K} - \mathbf{R}\mathbf{C}_E\mathbf{B}_E\mathbf{Q}]^{-1}\mathbf{R}$ matrix of a model with 260 segments with Dirichlet boundary conditions imposed. Notice that the diagonal still dominates the matrix but off diagonal terms are more significant than for the small system.




Figure 9. Performance of the monolithic and partitioned systems with 240 segments for two stapes boundary conditions: left hand side, zero Neumann (zero acceleration); and right hand side zero Dirichlet (zero pressure). The time step size for the monolithic and partitioned systems were $1 \mu\text{s}$ and $0.1 \mu\text{s}$ respectively.

are applied at the stapes the system admits a set of solutions. Figure 9 shows the behaviour of the monolithic and partitioned systems for the cases where the stapes boundary condition is zero Neumann (zero acceleration) and zero Dirichlet (zero pressure). The 240 segment model was adopted as the standard size model for this test and the rest of the experiments described. This choice is justified in the next section.

When Neumann boundary conditions were used, the energy norm of the partitioned system drifted away from the energy norm of the monolithic system as the simulation time increased, indicating that the partitioned system is slightly unstable for this case. The Dirichlet boundary condition seems to remedy this instability and the energy norm of both the partitioned system and the monolithic system both approach zero. It is uncommon for cochlear models to have only Neumann boundary conditions because the round window is a natural place to specify a zero pressure boundary condition. Therefore, we only consider the performance of systems using the Dirichlet boundary condition in the rest of the results.

7.3. Increasing the number of radial segments

240 segment models were used by [29] and [31], which are two papers that considered this cochlear model previously. It also provides a convenient size to adopt in this work. Those works were completed in the 1980s and computational constraints were an important consideration. Some modern cochlear models use significantly more segments. In general, the number of




Figure 10. Performance of the monolithic and partitioned system for various numbers of radial segments, as indicated on the figure. Dirichlet boundary conditions were used on the stapes. The trends for 16 and 240 segments were multiplied by four and offset by 5×10^{-17} and 2×10^{-16} respectively when plotting the figure. The time step size for the monolithic was $1 \mu\text{s}$ and for the partitioned systems a time step size of $0.1 \mu\text{s}$ was used.

segments used should be determined either by ensuring that increasing the number has negligible effects on the data produced (which those authors did as a matter of practice), or by anatomical considerations (there are about 3000 inner hair cells in man and this provides an upper limit on how small a radial segment should reasonably become). We test the effect of other choices for the number of segments in the this section.

Four simulations of different sizes were performed for the monolithic and partitioned systems. These were: 16, 60, 240 and 960 segments. We present the data in figure 10 using two panels. The first panel shows the results from the smaller two simulations and the second from the larger two simulations. Since the energy norm defined in equation (21) does not account for the size of each segment, we scale the data from the 16 and 240 segment simulations by a factor of 4 to ensure the values are of similar order of magnitude to the 60 and 960 simulations respectively. We also offset these trends by 5×10^{-17} and 2×10^{-16} respectively when plotting the figure.

The overall trend in the figures is that the performance of the partitioned system degrades as the number of segments used is increased. The slight improvement in performance between the 16 segment and the 60 segment model runs against this expected behaviour. We think it is due to the effect of the Dirichlet boundary condition on the structure of the fluid coupling matrix, which becomes increasingly significant as the matrix size is reduced.

7.4. Sensitivity to time step

The partitioned system evaluates the fluid model only once per time step compared to four times per time step in the monolithic system, if RK4 is used. This possible time saving can only be realised if the time step size required to obtain equivalent accuracy from the partitioned method is similar to that used in the monolithic method. In order to test this, we simulated the two systems for three (constant) time step sizes; in the monolithic case: $1 \times 10^{-5}\text{s}$, $1 \times 10^{-6}\text{s}$, and $1 \times 10^{-7}\text{s}$; and in the partitioned case $1 \times 10^{-6}\text{s}$, $1 \times 10^{-7}\text{s}$, and $1 \times 10^{-8}\text{s}$. The monolithic system produces similar results for the time steps of $1 \times 10^{-6}\text{s}$, and $1 \times 10^{-7}\text{s}$ suggesting that $1 \times 10^{-6}\text{s}$ is adequate to simulate the system. A time step of $1 \times 10^{-5}\text{s}$ is slightly too large as the trend deviates from the other results^{††}. The partitioned system is more sensitive to the time step size. A time step of $1 \times 10^{-6}\text{s}$ produces very poor results. The results for a time step of

^{††} $2.5 \times 10^{-6}\text{s}$ is a common choice of time step size in cochlear models. In the monolithic model it is limited by the fastest pole of the system, which can be found from the eigenvalues. We have not optimised the values for time step used in the rest of the paper.

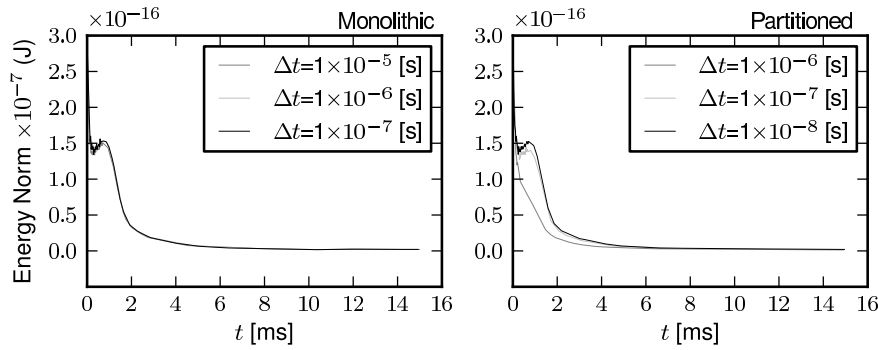


Figure 11. A figure showing the effect of time step size on the 240-segment-system's behaviour. Dirichlet boundary conditions were used at the stapes.




Figure 12. Behaviour of the 240-segment-system before ($R_{30} = -4.96$), in the vicinity of ($R_{30} = -4.9633$) and after ($R_{30} = -4.964$) the bifurcation point of the system ($R_{30} = -4.9634$). Dirichlet boundary conditions were used at the stapes.

1×10^{-7} s approach those for 1×10^{-8} s, but not as closely as the monolithic system's 1×10^{-6} s results and approach the 1×10^{-7} s results. Overall, to obtain similar levels of accuracy, the partitioned system needs to use a time step size that is 10 to 100 times smaller than that used in the monolithic system. This negates any time saving obtained by solving the fluid only once per time step.

These results verify that the partitioned system approaches the monolithic system when the time step size approaches zero, however they also suggest that the partitioned system is stiff compared to the monolithic system. A time step of 1×10^{-7} s was generally used for the partitioned system because of the poor performance for a time step of 1×10^{-6} s.

7.5. Sensitivity to R_{30} value

Finally, we investigated whether the partitioned method shows consistent performance as the bifurcation point of the system was crossed, as the results from the small system suggest it might. To do this, the bifurcation point of the system was found by studying the eigenvalues of the (linear) monolithic model as R_{30} was varied. The bifurcation was found to occur at $R_{30} = -4.9634$ to 4 decimal places. Values of R_{30} before (-4.96), in the vicinity of (-4.9633), and after (-4.964) bifurcation were selected and simulated. For a large system, the best performance was found before the bifurcation point. The result from the vicinity of the bifurcation is slightly degraded and the result after bifurcation shows further degradation. The selected values of R_{30} are closer to the bifurcation than those presented for the small system, but results from the small system that were not selected for presentation suggest that this is not the dominant factor.

8. DISCUSSION

8.1. Analysis of the reported instability

Section 5 investigated a report of instability in a common modeling approach. Analysis of the eigenvalues of the system confirmed the published bifurcation point for a 2 segment system. The stapes boundary conditions had a significant effect on the bifurcation point for the 2 segment model. This effect was not noticeable for the larger system – a finding that agrees with section 7.1. In that section it was found that the stapes boundary condition had a visible effect only on the fluid coupling of the radial segments in its vicinity. It is still conceivable that the stapes boundary condition could have a significant effect on the bifurcation of the whole system, if, for example, radial segments in the vicinity of the stapes were the closest to bifurcation.

Both analysis of eigenvalues and time domain simulation using the RK4 show that the FEM model is stable for $R_{30} = -4.96$, the value reported as unstable, but the bifurcation point for the FEM model is in the vicinity, at $R_{30} = -4.9634$. Given the proximity to the FEM model bifurcation point, it is possible that the Green's function model has a slightly different bifurcation point, meaning that the instability observed was correct behaviour. Matthews, who reported the instability, used a Green's function model extensively in his PhD thesis so the possibility that a software bug resulted in the instability seems unlikely.

The approximation $R_{3i} \approx -R_{2i}$ for when the second degree of freedom has a significant influence on the first degree of freedom only predicts the general vicinity of the bifurcation for the 240 segment system but is particularly accurate for the 2 segment, fixed stapes model. This contributed to the original expectation that $R_{30} = -4.96$ should be stable.

We have presented some time domain results that show that RK4 behaves as expected before, in the vicinity of and beyond the bifurcation point. Therefore, this method is suitable for this (type of) cochlear model in the vicinity of the critical point of the system.

The 2 segment model possibly provides an interesting view of “whole cochlear resonances” in action, as discussed in section 5.3.

8.2. Performance of the sequentially staggered partitioned system

In this paper we have attempted to find a sequentially staggered partitioned system that models a system known to have strong FSI coupling. We are interested in partitioned systems in general because they provide greater freedom to replace the fluid or structure models when desired. The deal.II cochlear modeling framework that we are developing (see section 5.2) already treats the radial segments as encapsulated C++ objects allowing them to be easily replaced by alternative segment models. Adding the ability to incorporate alternative fluid models seems to be a logical step towards multiscale models of the cochlea.

The decision to look for a sequentially staggered solution, despite the fact that strong coupling is present in the system, was partly motivated by a desire to investigate the simplest partitioned system before considering more advanced systems. It was also motivated by the knowledge that the current solution method of a Laplacian fluid model + RK4 time stepping is computationally cheap to implement and adequate for a wide range of situations. (This compensates for the fact that a large number of time steps are required to simulate a single physiologically relevant input, and numerous inputs must be simulated to characterise a model's performance.) If a partitioned method is to compete with the current solution method it must either offer similar computational requirements or a significant improvement in accuracy. Some cochlear models inherently require an implicit solver rather than the RK4. ([43] and [33] are examples.) For these systems, an iteratively staggered partitioned system would be a natural choice.

The sequentially staggered method was based on separating the simulation into fast and slow feedback paths. Simplified models of the fast feedback paths were provided to both the fluid and the structure separately. The segment model considered naturally separates into these components. The Laplacian fluid model, however, does not. For a small system, it separates into

local and nonlocal components, with the contribution of the nonlocal components significantly smaller than that of the local components. While this pattern extends to larger systems, the neighbouring segments have a significant effect. This violates one of the assumptions of the partitioning.

The promising behaviour of the small system did not carry over to the larger system. Although the property that the partitioned system converges to the monolithic system as the time step size is reduced was observed, a step size between a factor of 10 and 100 times smaller was required to obtain similar accuracy to the monolithic solution. The partitioned method solves the fluid model once per time step vs 4 times per time step for the monolithic method using RK4, but this potential speed up is completely offset by the small time step size required. This would defeat the point of using this partitioned method for a Laplacian fluid model. Therefore, in this form the sequentially staggered partitioned method is a failure.

The assumption that is being violated is the assumption that segment self coupling dominates the fluid feedback (equation 27). Therefore, the accuracy of the sequentially staggered partitioned method could be improved by keeping $2n$ of the neighbouring minor diagonals^{‡‡}:

$$\mathbf{G}_{jk} = \begin{cases} (\mathbf{Q}[\mathbf{K} - \mathbf{R}\mathbf{C}_E\mathbf{B}_E\mathbf{Q}]^{-1}\mathbf{R})_{jk} & |j - k| \leq n \\ 0 & \text{otherwise} \end{cases}. \quad (37)$$

This would cause coupling terms between individual segments. Some microscopic models of the cochlea already incorporate so-called feedforward coupling ([44] and [43] are examples) so a similar method could be used to incorporate the coupling terms arising due to the fluid. Note that as n is increased the system approaches the monolithic system.

8.3. Other fluid models

As we noted above, the desire to offer alternatives to the Laplacian fluid model was a major reason for this investigation into partitioned cochlear models. As we mentioned in section 3, it is known that including viscous behaviour in a cochlear model will improve the low frequency response, and compressibility the high frequency response. The errors introduced by neglecting these effects are generally small, but as our understanding of the cochlea develops and the computational expense of including them becomes less prohibitive, it seems logical to expect that they will be more widely used. However, if a partitioned approach is used to include these small effects, the error associated with the method itself must be smaller still.

Whilst the Laplacian fluid model does not contain time dependent terms, the Navier-Stokes equations, which are widely used to model viscous fluids, do. In fact, the Laplace equation is a high frequency model of the linear Navier-Stokes equations. Therefore, the sequentially staggered partitioned method describe in this paper could be tested with the linear Navier-Stokes equations used once per time step to calculate $\mathbf{P}(t)$ and the Laplace equation used to compute $\mathbf{P}_f(t)$. The fact that the Navier-Stokes equations require velocities whereas the Laplace equation requires accelerations would need to be taken into account. Similarly other implementation details arise because the Navier-Stokes equations lead to a vector system whilst the Laplace equation leads to a scalar system. However, since the Laplace equation is a good model of the fast behaviour of the fluid model, we think that this is a promising combination.

Including compressibility in a fluid reduces the speed of sound in the fluid from a theoretical infinity in the incompressible case to some smaller amount. In a cochlear model, this means

^{‡‡}Keeping just a few minor diagonals can have significant effect on the system's performance. For example, when $n = 3$ for the 240-segment model \mathbf{G}_{jk} accounts for over 95% of $\mathbf{Q}[\mathbf{K} - \mathbf{R}\mathbf{C}_E\mathbf{B}_E\mathbf{Q}]^{-1}\mathbf{R}$ (based on the size of terms included and excluded). The error between the energy norm of the partitioned method using this \mathbf{G}_{jk} and the monolithic system is on the order of 1×10^{-18} for a step size of 1×10^{-6} s in both systems. For comparison, the error for the partitioned system with $n = 0$ is on the order of 1×10^{-17} for a step size of 1×10^{-7} s (see figure 11). This shows the major improvement just 6 minor diagonals have made.

that the fast coupling between remote segments is reduced which would appear to be beneficial for the partitioned method if anything because feedback from the local area of the segment is emphasised.

9. CONCLUSIONS

In this paper we have investigated an instability in a popular (monolithic) cochlear modeling approach described in the literature. The eigenvalues of the model described as unstable were computed and found to be very close to the bifurcation point expected for the FEM implementation of the model. It is reasonable that differences between the Green's function implementation and the FEM implementation could mean that the two systems lie on opposite sides of the bifurcation. Comparison of time domain simulations and computed eigenvalues of the model showed that both the 240-segment simulations and the small 2-segment simulations (originally used to investigate the instability by the authors who noticed it) behave as expected in the vicinity of the bifurcation. Therefore, we conclude that the described instability is not a cause for concern for users of the monolithic model. We followed up this work by investigating a method for a sequentially staggered partitioned cochlear model using a paradigm of separating the computation into "fast" and "slow" feedback components. This is seen as a natural extension to current practice. We were aware that the cochlea is known to have strong FSI coupling that make this approach difficult, but wanted to test the simple alternatives before the complex. We were also conscious that the monolithic solution using the Laplace equation is computationally cheap to solve and any partitioned method must maintain this property to be adopted. The results from the small system were mixed, and when extended to larger systems certain assumptions did not hold. A discussion of where the assumptions are violated, and some alternatives to avoid this, is presented. This first approach fails because the assumptions factoring the fluid model into "fast" and "slow" components fails. Given that the Laplace equation is essentially the "fast" component of the linear Navier-Stokes equation, the general paradigm may still be useful for generating partitioned time stepping methods for cochlear models. We hope that a method like this might one day find its place in multiscale cochlear models, because a consideration of the cochlear anatomy, and a review of the methods that have been applied to model the cochlea, suggest that multiscale cochlear models will be an increasingly useful compliment to ongoing cochlear research.

ACKNOWLEDGEMENT

We would like to thank the organisers of AfriCOMP 2011 for an engaging conference, and Dr. McBride and Prof. Steinmann for an interesting mini symposium at that conference. One of us (Michael Rapson) would like to thank the deal.II users group in CERECAM for the many fruitful discussions on topics relevant to this paper. The financial assistance of the National Research Foundation (NRF, South Africa) towards this research is hereby acknowledged. Opinions and conclusions arrived at, are those of the authors and are not necessarily to be attributed to the NRF.

References

1. Passchier-Vermeer W, Passchier WF. Noise exposure and public health. *Environ Health Perspect* Mar 2000; **108 Suppl 1**:123–131.
2. Allen J. *Physiology of the Ear*, chap. Nonlinear Cochlear Signal Processing. Second edition edn., Singular Thomson Learning; San Diego, CA, 2001; 393–442.
3. Lyon RF, Mead C. An analog electronic cochlea. *IEEE Trans-ASSP* Jul 1988; **36**:1119–1134, doi:10.1109/29.1639.
4. Pickles JO. *An Introduction to the Physiology of Hearing*. 2nd edn., Academic Press, 1988.
5. Dallos P, Popper AN, Fay RR (eds.). *The Cochlea*. Springer Handbook of Auditory Research, Springer-Verlag New York, Inc., 1996.
6. Thorne M, Salt AN, DeMott JE, Henson MM, Henson OW, Gewalt SL. Cochlear fluid space dimensions for six species derived from reconstructions of three-dimensional magnetic resonance images. *The Laryngoscope* October 1999; **109**:1661–1668, doi:10.1097/00005537-199910000-00021.

7. Kessel RG, Kardon RH. *Tissues and Organs: a text-atlas of scanning electron microscopy*. W. H. Freeman, San Francisco, 1979.
8. von Békésy G. *Experiments in Hearing*. McGraw-Hill, New York, 1960.
9. Matthews JW. Mechanical modeling of nonlinear phenomena observed in the peripheral auditory system. PhD Thesis, Washington University, St. Louis, MO 1980.
10. Peterson LC, Bogert BP. A dynamical theory of the cochlea. *J. Acoust. Soc. Am.* May 1950; **22**(3):369–381, doi:10.1121/1.1917149.
11. Ranke OF. Theory of operation of the cochlea: A contribution to the hydrodynamics of the cochlea. *J. Acoust. Soc. Am.* 3 Aug 1950; **22**(6):772–777, doi:10.1121/1.1906688.
12. Fletcher H. On the dynamics of the cochlea. *J. Acoust. Soc. Am.* Nov 1951; **23**(6):637–645, doi:10.1121/1.1906813.
13. Allen JB, Sondhi MM. Cochlear macromechanics: Time domain solutions. *J. Acoust. Soc. Am.* Jul 1979; **66**(1):123–132, doi:10.1121/1.383064.
14. Cai H, Manoussaki D, Chadwick R. Effects of coiling on the micromechanics of the mammalian cochlea. *J. R. Soc. Interface* 16 Jun 2005; **2**:341–348, doi:10.1098/rstb.2005.0049.
15. Kolston PJ, Ashmore JF. Finite element micromechanical modeling of the cochlea in three dimensions. *J. Acoust. Soc. Am.* Jan 1996; **99**(1):455–467, doi:10.1121/1.414557.
16. Neely ST, Kim DO. A model for active elements in cochlear biomechanics. *J. Acoust. Soc. Am.* May 1986; **79**(5):1472–1480, doi:10.1121/1.393674.
17. de Boer E. Auditory Physics. Physical Principles in Hearing Theory. III. *Phys. Repts.* 1991; **203**(3):125–231, doi:0370-1573(84)90108-X.
18. Parthasarathi AA, Grosh K, Nuttall AL. Three-dimensional numerical modeling for global cochlear dynamics. *J. Acoust. Soc. Am.* Jan 2000; **107**(1):474–485, doi:10.1121/1.428352.
19. Lim KM, Steele CR. A three-dimensional nonlinear active cochlear model analyzed by the wkb-numeric method. *Hear Res* Aug 2002; **170**(1-2):190–205, doi:S0378-5955(02)00491-4.
20. Neely ST. Finite difference solution of a two-dimensional mathematical model of the cochlea. *J. Acoust. Soc. Am.* May 1981; **69**(5):1386–1393, doi:10.1121/1.385820.
21. Allen JB. Two-dimensional cochlear fluid model: New results. *J. Acoust. Soc. Am* Jan 1977; **61**(1):110–119, doi:10.1121/1.381272.
22. Lüling H, Franosch JMP, van Hemmen JL. A two-dimensional cochlear fluid model based on conformal mapping. *J Acoust Soc Am* Dec 2010; **128**(6):3577–3584, doi:10.1121/1.3505108.
23. Puria S, Allen JB. A parametric study of cochlear input impedance. *J. Acoust. Soc. Am.* Jan 1991; **89**(1):287–309, doi:10.1121/1.400675.
24. Neely ST, Allen JB. *Proceedings of the 10th Mechanics of Hearing Workshop*, chap. Retrograde waves in the cochlea. World Scientific Publishing Co., 2009; 62–67.
25. Vieregger MA, de Boer E. Matching impedance of a nonuniform transmission line: Application to cochlear modeling. *J. Acoust. Soc. Am.* Jan 1987; **81**(1):184–186, doi:10.1121/1.395135.
26. Vieregger M. Mechanics of the inner ear. PhD Thesis, Delft University of Technology 2 Oct 1980.
27. Ramamoorthy S, Zha DJ, Nuttall AL. The biophysical origin of traveling-wave dispersion in the cochlea. *Biophys J* Sep 2010; **99**(6):1687–1695, doi:10.1016/j.bpj.2010.07.004.
28. Givberg E, Bunn J. A comprehensive three-dimensional model of the cochlea. *J. Comp. Phys.* 11 2003; **191**(2):377–391, doi:10.1016/S0021-9991(03)00319-X.
29. Neely ST. Fourth-order partition dynamics for a two-dimensional model of the cochlea. PhD Thesis, Washington University, Saint Louis, Missouri August 1981.
30. Elliott SJ, Ku EM, Lineton B. A state space model for cochlear mechanics. *J. Acoust. Soc. Am.* Nov 2007; **122**(5):2759–2771, doi:10.1121/1.2783125.
31. Matthews JW, Molnar CE. On the stability of cochlear mechanical models. *Cochlear Mechanisms - Structure, Function and Models*, Wilson JP, Kemp DT (eds.). Plenum Press, 1989; 437–444.
32. Diependaal RJ, Duifhuis H, Hoogstraten HW, Vieregger MA. Numerical methods for solving one-dimensional cochlear models in the time domain. *J. Acoust. Soc. Am.* Nov 1987; **82**(5):1655–1666, doi:10.1121/1.395157.
33. Moleti A, Paternoster N, Bertaccini D, Sisto R, Sanjust F. Otoacoustic emissions in time-domain solutions of nonlinear non-local cochlear models. *J Acoust Soc Am* Nov 2009; **126**(5):2425–2436, doi:10.1121/1.3224762.
34. Duifhuis H. Hopf-bifurcations and Van der Pol oscillator model of the mammalian cochlea. *What Fire is in Mine Ears: Progress in Auditory Biomechanics*, American Institute of Physics: Melville, NY, 2011; 195–201.
35. Hamilton TJ, Tapson J, Rapson M, Jin C, van Schaik A. Understanding the mathematics of hearing using electronic circuits. *ANZIAM J.* 2010; **51**:C300–C315.
36. Bangerth W, Hartmann R, Kansch G. deal.II — a general-purpose object-oriented finite element library. *ACM Trans. Math. Softw.* Aug 2007; **33**(4), doi:10.1145/1268776.1268779.
37. Bangerth W, Kansch G. *deal.II Differential Equations Analysis Library, Technical Reference*. URL <http://www.dealii.org>, last seen: 14/05/2011.
38. Jones E, Oliphant T, Peterson P, et al. SciPy: Open source scientific tools for Python. <http://www.scipy.org>, last seen 22/10/2010 2001–.
39. Iserles A. *A First Course in the Numerical Analysis of Differential Equations*. Second edition edn., Cambridge Texts in Applied Mathematics, Cambridge University Press, 2009.
40. Neely ST. Mathematical models of the mechanics of the cochlea. Master's Thesis, California Institute of Technology, Pasadena, California 1978.
41. Küttler U, Gee M, Förster C, Comerford A, Wall WA. Coupling strategies for biomedical fluid-structure interaction problems. *Int. J. Numer. Meth. Biomed. Engng.* 2010; **26**:305–321, doi:10.1002/cnm.1281.

42. Förster C, Wall WA, Ramm E. Artificial added mass instabilities in sequential staggered coupling of nonlinear structures and incompressible viscous flows. *Comput. Methods Appl. Mech. Engrg.* 2007; **196**:1278–1293, doi:10.1016/j.cma.2006.09.002.
43. Kim Y, Xin J. A two-dimensional nonlinear nonlocal feed-forward cochlear model and time domain computation of multitone interactions. *Multiscale Model. Simul.* 2005; **4**(2):664–690.
44. Geisler CD, Sang C. A cochlear model using feed-forward outer-hair-cell forces. *Hearing Res.* 21 Mar 1995; **86**(1,2):132–146, doi:10.1016/0378-5955(95)00064-B.

Appendix E

Capturing the Low Frequency Cochlear Impedance in Time Domain Models: The Role of Viscosity

E.1 Introduction

The following presentation was given at the 159th meeting of the Acoustical Society of America (ASA) in Baltimore, 2010 ([Rapson and Tapson 2010](#)). Preparing the presentation highlighted to me the difficulty of accurately capturing the low frequency response of the cochlea, given that fluid viscosity plays a major role in determining that response. It had a strong influence on the discussion in sections 4.2.5 and 4.6.2.

The work compared two simulations, solving the Navier-Stokes equations and potential fluid formulation, respectively. The major problem with simply solving the Navier-Stokes equations is that they are inherently computationally intensive, and capturing the viscous boundary layer requires a high mesh refinement. Two separate methods that aimed to incorporate the effects of the fluid viscosity into the potential fluid formulation were then tested. The first of these methods was the widely used technique of lumping the viscous effects into the helicotrema, which is modelled separately. The second method tested was an experimental approach that aimed to represent the distributed effects of the viscosity. The promising initial results from the second method required an increasingly complex implementation. Work on the method ultimately stopped short of either a final implementation or an obvious failure. That is, the method failed in the mode of “Iterative Improvements Become Increasingly Prohibitive” discussed in section 7.1.2.2.

I include this work as an appendix for three reasons: to ensure that the work associated with the ASA meeting abstract is available to interested parties, as it has not been published elsewhere; to demonstrate some of the difficulties associated with handling viscosity in time domain computations,

which justifies the various compromises used in models; and to give an example of the failure mode discussed in section 7.1.2.2. The remainder of the section presents the article abstract submitted, and the slides and script presented at the meeting.

E.2 Original Article Abstract


Viscous effects in the cochlear fluid have been shown to be negligible for frequencies greater than 100 Hz M. Viergever, PhD thesis. However they become important at lower frequencies, hence viscous damping is sometimes included in helicotrema models [S. T. Neely and D. O. Kim, J. Acoust. Soc. Am. 70, 1172-1480 (1986). Modeling viscosity in time-domain computational cochlear models for wide-band input signals is problematic. Since the boundary layer thickness is inversely proportional to frequency, Navier-Stokes formulations require fine spatial resolutions, and potential fluid formulations are inherently inviscid models. Lumping all viscous effects into the helicotrema is not physiologically accurate as 'the helicotrema does not appear to be a dominant constriction' [Lynch *et al.*, J. Acoust. Soc. Am. 72, 108-130 (1982)]. An alternative is proposed whereby a potential fluid formulation is augmented with frequency dependent viscosity corrections that are based on prior analysis of the cochlear geometry and become negligible at higher frequencies. For simplified fluid boundary conditions, the predictions of the model are compared to Navier-Stokes, potential, and lumped viscosity models for frequencies in the range 10 Hz to 10 kHz. The simulations are also compared to physiological data for the input impedance of the cochlea.

E.3 Slides Presented

<p>Capturing low frequency cochlear impedance in time domain models: the role of viscosity</p> <p>By Michael J. Rapson and Jonathan C. Tapson Dept. of Electrical Engineering University of Cape Town, South Africa</p>	<p>Low frequencies (< 500 Hz) and time simulation</p> <p>Problematic low frequency behavior:</p> <ul style="list-style-type: none">• reflections at apical basilar membrane boundary• whole cochlear resonances• [Neely MSc] and [Matthews Dsc] <p>Possible causes discussed in the literature:</p> <ul style="list-style-type: none">• incorrect helicotrema impedance• neglecting viscosity• neglecting scalae tapering• inaccurate cochlear frequency-place map• Dalos [1970], Lynch <i>et al.</i> [1982] Puria and Allen [1991]
--	---

Methodology


- 4 FEM formulations were implemented
- Boundary conditions taken from input impedance network by Lynch et al. (Right)
- FEM simulations compared to each other and the network predictions



Outline

- Obtaining the cochlear geometry
- Benchmark fluid formulations
 - Potential model and Navier Stokes model
 - Typical flow and pressure profiles
 - Comparison with Lynch et al. Network
- Corrections to the potential model
 - Lumped viscosity profiles
 - Outline of proposed distributed viscosity model
- Conclusions


Cochlear Geometry



Data from Mountain et al. (2003) & Sato et al. (1999)

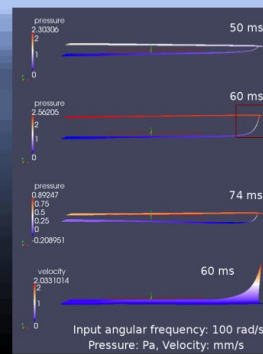
Scalae areas converted to heights by $h \approx \sqrt{A}$

- Preserves heights
- Area not preserved



Navier-Stokes Formulation

- Adapted from code by Martin Kronbichler
- Models the fluid viscosity
- Is computationally intensive to run
- Pressure profile is time dependent
- Shows a major pressure drop in apical region of scala tympani




Standard Potential Formulation


- Solves the equation $\nabla^2 \phi = 0$ for a potential field, ϕ
- Pressure and velocity post-processed using:

$$\mathbf{v} = -\nabla \phi \quad p = \frac{\partial}{\partial t} \rho \phi$$

- Assumes fluid is inviscid
- Potential profile shown above is typical of all time instants
- Pressure profiles follow the same trend



Comparison to Network



Red: Pressure [Pa], Blue: Potential [mm²s⁻¹]. Solid lines are predictions from detailed network. Filled shapes are single frequency simulations. Dashed line are from a step test.

Lumped Damping

$$v = \frac{s\rho}{sM+R} \Delta(A\phi)$$

- Models a part of the domain using a transfer function, with viscous damping
- Potential profile is time dependent, but shows a step change not present in Navier-Stokes results

Distributed Damping

$$\nabla \cdot \mathbf{v} = 0, p = s\rho\phi \quad (1)$$

$$\mathbf{v} = - \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}^T D(x, y, t) - \nabla\phi \quad (2)$$

$$\nabla^2\phi + \frac{\partial}{\partial x} D(x, y, t) = 0 \quad (3)$$

$$sD(x, y) = \frac{\mu}{\rho} w(s)g(x, y)^2 v_x \quad (4)$$

$$\partial^2 v_x / \partial y^2 \approx -g(x, y)^2 v_x \quad (5)$$

These apply in the marked region:

Distributed Damping

- Scalae tapering and an integration constant must be accounted for in velocity equation
- The modifications must become negligible at high frequencies
- Preliminary results suggest these can be achieved

Conclusions

Accurately modeling viscosity at low frequencies is problematic but:

- Distributed damping approaches may be an option
- Scalae tapering may reduce the resonance peak in potential models, reducing apical reflections (supports Puria and Allen [1991])
- Using a frequency map that goes to zero as frequency approaches zero avoids reflections (Puria and Allen [1991])

When scalae tapering is present, converting cross-sectional areas to lengths for 2D finite element models distorts volume flow rates

Acknowledgments

The financial assistance of the National Research Foundation (NRF, South Africa) towards this research is hereby acknowledged. Opinions and conclusions arrived at, are those of the author and are not necessarily to be attributed to the NRF.

Michael would like to thank his supervisor, Jonathan Tapson and his research group, as well as CERECAM and the UCT.deal.II users group.

deal.II – a finite element Differential Equations Analysis Library

ParaView – data analysis and visualization

References

Neely, S. T., "Mathematical models of the mechanics of the cochlea", MSc thesis, *California Institute of Technology, Pasadena, California*, 1978

Matthews, J. W., "Mechanical modeling of nonlinear phenomena observed in the peripheral auditory system", PhD thesis, *Washington University, St. Louis, MO*, 1980

Dallos, P., "Low-Frequency Auditory Characteristics: Species Dependence", *J. Acoust. Soc. Am.*, 1970, 48, 489-499

T. J. Lynch, I. Nedzelitsky, V. & Peake, W. T., "Input impedance of the cochlea in cat", *J. Acoust. Soc. Am.*, 1982, 72, 108-130

Puria, S. & Allen, J. B., "A parametric study of cochlear input impedance", *J. Acoust. Soc. Am.*, 1991, 89, 287-309

Mountain, D. C., Hubbard, A. E., Ketten, D. R., & O'Malley, J. T., "The helicotrema: Measurements and models" in *Biophysics of the Cochlea: From Molecules to Models*, Gummer, A. W., Dalhoff, E., Nowotny, M., & Scherer, M. P. (ed.), World Scientific, Singapore, 2003

Sato, M., Leake, P. A., & Hradek, G. T., "Postnatal development of the organ of Corti in cats: a light microscopic morphometric study", *Hearing Res.*, 1999, 127, 1-13

E.4 Presentation Script

(Slide One) Good morning, my name is Michael Rapson and I am pursuing a PhD in electrical engineering at the University of Cape Town in South Africa. Jonathan Tapson is my supervisor, and in this presentation I would like to discuss some of our work on understanding the low frequency behaviour of time domain cochlear simulations. I will refer to a number of other studies that directly or indirectly address this topic, and this work agrees with them for the most part. The contribution

in this work is to consider the application to finite element models of the cochlea, in particular two dimensional models.

(Slide Two) At low frequencies standing waves, otherwise described as whole cochlear resonances, have been observed and attributed to reflections at the apical basilar membrane boundary. In frequency domain models the problem frequencies can be avoided, but not in time domain models. A number of possible causes have been discussed in the literature. These are: including incorrect helicotrema impedance, neglecting viscosity and scalae tapering, or using an inaccurate cochlear map. Of these, viscosity and tapering are closely related. It has been widely assumed that the helicotrema impedance has a significant viscous damping action, but in fact none of the references cited below support this.

(Slide Three) In this work 4 finite element fluid formulations are simulated with boundary conditions implementing similar stapes and round window models. The boundary conditions were based on the published input impedance network. To focus on the fluid flow through the helicotrema, the organ of Corti complex was neglected. It is modelled as infinitely stiff, effectively removing R_c in the diagram. The finite element simulation results were run in the time domain for specific input frequencies, and the simulation results were compared to the network predictions.

(Slide Four) Having mentioned the problems associated with modelling low frequencies, and discussed the methodology used, I will speak briefly on each of the fluid formulations. I will use a qualitative approach to the data, mainly presenting typical flow and pressure profiles. Trends in the data will be compared against the detailed network. Two corrections to the potential model that account for viscosity will be presented: a method where all the viscous effects are modelled in a small region; and the proposed distributed viscosity model. Finally conclusions are drawn.

(Slide Five) Additional data was required for the cochlear geometry. The published cross-sectional area data was converted to heights for a two dimensional model by taking the square root. This preserves the height of the scalae, which is important for viscous fluid models, because the ratio of the height to boundary layer thickness is not changed. However a two dimensional model implicitly assumes a constant scalae width, so the implicit cross-sectional area is not preserved¹. The scalae data for the cat shows significant narrowing in the apical region of the scala tympani, highlighted in red. It must be noted that other species do not show such a radical narrowing.

(Slide Six) The first fluid formulation used was an implementation of the Navier-Stokes equations by Martin Kronbichler. His code was adapted to include stapes and round window boundary conditions. Navier-Stokes equations model fluid viscosity, hence these are taken as the accurate solution. However, they are computationally expensive to solve. Furthermore, if a wide frequency range is modelled, a small time step and high mesh refinement to capture the viscous boundary layer in the fluid flow are required. On the right, I show the pressure profile at three time instants. Notice that

¹This can be corrected for by allowing the density of the fluid in the model to vary.

the major area of pressure drop is in the apical region of the scala tympani. The velocity profile has a similar shape at all times, however the size varies. A high pressure drop in the scala tympani is associated with high flows.

(Slide Seven) Solving a Laplace equation for either a pressure or potential field is much less computationally intensive. In the potential case, the velocity and pressure are post-processed using the relationships shown. This formulation assumes an inviscid fluid, and one of the effects of this is that the potential, and hence the pressure profile, is similar at all time instants. Therefore, at these low frequencies the standard potential formulation gives qualitatively different results to the Navier-Stokes formulation. At higher frequencies the solutions converge.

(Slide Eight) The input impedance of the two formulations were compared against each other and against bode plots of detailed network with and without viscous damping. All inviscid results are input potentials, shown in blue, and results with viscosity are input pressures, shown in red. Results from the network show a resonance peak, and the inviscid peak is about 6 dB higher than the viscous peak. This peak occurs in the range of frequencies that whole cochlear resonances were reported at. The potential model does not show a resonance peak, which would seem to be due to the scalae tapering. If this is the case, it would suggest that the scala tympani tapers too much in this geometry. The network and simulation results both show that at low frequencies, the pressure in viscous models is in phase with the potential in inviscid model. In fact, they should be out of phase by ninety degrees, like they are at high frequencies. Due to numerical problems, the Navier-Stokes model could only be solved below the resonance peak.

(Slide Nine) A common approach to correct the low frequency behaviour is to model a part of the fluid domain using a transfer function with damping. All the effects of fluid viscosity are assumed to act in this region, which is usually the helicotrema, A in the insert. In the cat this is not justified since the scala tympani is much narrower than the helicotrema. Hence I modelled the viscous effects in region B. In this implementation, some mass was also required for causality of the transfer function. The value of R was set equal to the damping in the detailed network, which turned out to be too large, hence I present results for a small time range only. Overall, the changes have a positive effect on the phase of the input potential and the potential profile shows time dependence, however there is a step change in potential rather than the smooth change seen in the Navier-Stokes results.

(Slide Ten) We have investigated various ways to incorporate some viscous effects into the fluid model in a distributed manner. The most promising is inspired by the form of the Navier-Stokes equations: Starting from the continuity equation, and the desired relationship between pressure and potential, an expression for velocity is proposed which is equation (2). This leads to a differential equation to solve on the fluid domain, (3). $D(x,y,t)$ is defined by a transfer function, that is related to the Navier-Stokes equations, because the computational results show that one of the viscous terms dominates, and can be approximated well by equation 5, where the spatial function $g(x,y)$ will need to be fitted for each geometry. To illustrate this, plots of the dominant second derivative and its

estimate with a simple $g(x,y)$ are given. These equations apply in the marked region where viscous effects are large, the normal equations apply elsewhere.

(Slide Eleven) There are some further subtleties: the velocity equation must be modified to account for an integration constant and scalae tapering. Corrections must also be applied at the boundaries of the two fluid regions. Finally, the modifications must become negligible at high frequencies. Experiments with corrections fitted by hand suggest that these requirements can be achieved. Panel (a) shows the estimate of velocity in a viscous model showing close agreement with the actual velocity. (b) and (c) have a constant $D(x,y)$ term, but varying input velocities. As expected when input velocity is small the viscosity has a larger comparative effect. Also notice that it steepens the gradient in the apical region of the scala tympani. (d) shows that at high frequency the distributed damping becomes similar to the standard model. To move forward from here, the corrections must be calculated in the simulation, rather than by hand, and the method must be tested with a basilar membrane present. However, we are cautiously optimistic.

(Slide Twelve) This work has shown that accurately modelling viscosity at low frequencies in finite element models remains problematic. The proposed distributed damping formulation needs further work, but it or similar approaches may be an option in the future. It was observed that scalae tapering may reduce the resonance peak in models, reducing the problematic apical reflections. This supports previous findings. Another important result in that paper is that apical reflections can be eliminated completely by using certain cochlear frequency maps. While these maps do not match the best neural data, this approach will be useful for a wide range of experiments when the low frequency behaviour is not critical. When scalae tapering is present, a choice needs to be made between preserving the linear distances in the cochlear geometry, or preserving the volume flow rates. For the models solving the Laplace equation, preserving volume flow rates may prove more important.

(Slide Thirteen) I would like to invite any questions.

(Slide Fourteen – displayed while fielding questions.) Jont Allen raised the question of whether I had incorporated the nonlinear advective terms in the Navier-Stokes simulation, and if I had compared the nonlinear Navier-Stokes to linear Navier-Stokes code to assess the importance of the advective term. He noted that previous investigations suggest the nonlinear terms are negligible. I responded that the Navier-Stokes code did include nonlinearities, and I had not attempted to compare its results to a code that did not include nonlinearities. However, the nonlinear solver always completed in the minimum number of linear iterations, which strongly implies that a linear code is adequate. I agreed that the literature indicated the nonlinear terms are negligible, and referred the dimensional analysis performed by Max Viergever that was reviewed in section 4.2.2. The advective terms are functions of intensity, so care was taken to use realistic input sizes in the tests to ensure that the nonlinear advective terms were correctly scaled.

Bibliography

- Allen, J. (2001). *Physiology of the Ear*, chapter Nonlinear Cochlear Signal Processing, 393–442, second edition edition (Singular Thomson Learning: San Diego, CA).
- Allen, J. (2008). *Springer Handbook of Speech Processing*, chapter Nonlinear cochlear signal processing and masking in speech perception, 27–62, Springer Handbook of Series (Springer).
- Allen, J. B. (1977). “Two-dimensional cochlear fluid model: New results”, *J. Acoust. Soc. Am* **61**, 110–119.
- Allen, J. B. (2009). “Time-domain impedance”, Latest version: <http://auditorymodels.org/tmp/IEEE-Impedance.09/Impedance.pdf>.
- Allen, J. B. and Sondhi, M. M. (1979). “Cochlear macromechanics: Time domain solutions”, *J. Acoust. Soc. Am.* **66**, 123–132.
- Ashmore, J. F. (1983). “Frequency tuning in a frog vestibular organ”, *Nature* **304**, 536–538.
- Babbs, C. F. (2011). “Quantitative reappraisal of the Helmholtz-Guyton resonance theory of frequency tuning in the cochlea”, *J. of Biophys.*
- Balay, S., Brown, J., Buschelman, K., Gropp, W. D., Kaushik, D., Knepley, M. G., McInnes, L. C., Smith, B. F., and Zhang, H. (2011). “PETSc Web page”, <HTTP://www.mcs.anl.gov/petsc>.
- Bangerth, W., Burstedde, C., Heister, T., and Kronbichler, M. (2010). “Algorithms and data structures for massively parallel generic adaptive finite element codes”, Submitted to *ACM Transactions on Mathematical Software*.
- Bangerth, W., Hartmann, R., and Kanschat, G. (2007). “deal.II — a general-purpose object-oriented finite element library”, *ACM Trans. Math. Softw.* **33**, 24/1–24/27.
- Bangerth, W. and Kanschat, G. (2011). *deal.II Differential Equations Analysis Library, Technical Reference*, URL <http://www.dealii.org>, last seen: 18/01/2012.
- Bell, A. (2008). “The pipe and the pinwheel: Is pressure an effective stimulus for the 9 + 0 primary cilium?”, *Cell Biology International* **32**, 462–468.

- Bell, A. (2010). “The cochlea as a graded bank of independent, simultaneously excited resonators: Calculated properties of an apparent ‘travelling wave’”, in *Proceedings of 20th International Congress on Acoustics, ICA 2010*.
- Bell, J. A. (2005). “The underwater piano: A resonance theory of cochlear mechanics”, Ph.D. thesis, The Australian National University.
- Bogert, B. P. (1951). “Determination of the effects of dissipation in the cochlear partition by means of a network representing the basilar membrane”, *J. Acoust. Soc. Am.* **23**, 151–154.
- Bouzeghoub, M., Gardarin, G., and Valduriez, P. (1997). *Object Technology: Concepts and Methods* (International Thomson Computer Press), translated by Jack Howlett.
- Brandenburg, K. (1999). “MP3 and ACC explained”, AES 17th International Conference on High Quality Audio Coding.
- Braun, M. (1994). “Tuned hair cells for hearing, but tuned basilar membrane for overload protection: evidence from dolphins, bats, and desert rodents”, *Hear. Res.* **78**, 98–114.
- Braun, M. (2009). *Concepts and Challenges in the Biophysics of Hearing*, chapter Dual tuning in the mammalian cochlea: dissociation of neural and basilar membrane responses at supra-threshold sound levels – a meta-analysis, 162–167 (World Scientific, Singapore).
- Brownell, W. E., Bader, C. R., Bertrand, D., and de Ribaupierre, Y. (1985). “Evoked mechanical responses of isolated cochlear outer hair cells”, *Science* **227**, 194–196.
- Butcher, J. and Wanner, G. (1996). “Runge-kutta methods: some historical notes”, *Appl. Num. Math.* **22**, 113–151.
- Cai, H., Manoussaki, D., and Chadwick, R. (2005). “Effects of coiling on the micromechanics of the mammalian cochlea”, *J. R. Soc. Interface* **2**, 341–348.
- Caplan, A. L. (1983). “Beastly conduct: ethical issues in animal experimentation.”, *Ann. N. Y. Acad. Sci.* **406**, 159–169.
- CIOMS (1985). “CIOMS international guiding principles for biomedical research involving animals.”, *Altern. Lab. Anim.* **12**, ii–ii.
- Coddington, E. A. and Levinson, N. (1955). *Theory of Ordinary Differential Equations* (McGraw-Hill, New York).
- Cole, J. D. and Chadwick, R. S. (1977). “An approach to mechanics of the cochlea”, *J. of Appl. Math. and Phys. (ZAMP)* **28**, 785–804.

- Crawford, A. C. and Fettiplace, R. (1981). “An electrical tuning mechanism in turtle cochlear hair cells”, *J. Physiol.* **312**, 377–412.
- Dahl, O. J. and Nygaard, K. (1966). “SIMULA: An ALGOL-based simulation language”, *Communications of the ACM* **9**, 671–678.
- Dallos, P. (1970). “Low-frequency auditory characteristics: Species dependence”, *J. Acoust. Soc. Am.* **48**, 489–499.
- Dallos, P. (1996). *The Cochlea*, chapter Overview: Cochlear NeuroBiology, 1–43, Springer Handbook of Auditory Research (Springer-Verlag, New York).
- Dallos, P., Popper, A. N., and Fay, R. R., eds. (1996). *The Cochlea*, Springer Handbook of Auditory Research (Springer-Verlag New York, Inc.).
- Davis, T. A. (2004). “Algorithm 832: Umfpack v4.3—an unsymmetric-pattern multifrontal method”, *ACM Transactions on Mathematical Software* **30**, 196–199.
- Dawkins, R. (1986). *The Blind Watchmaker* (Longman Scientific & Technical).
- de Boer, E. (1980). “Auditory physics. Physical principles in hearing theory. I”, *Phys. Repts.* **62**, 87–174.
- de Boer, E. (1984). “Auditory physics. Physical principles in hearing theory. II”, *Phys. Repts.* **105**, 141–226.
- de Boer, E. (1991). “Auditory physics. Physical principles in hearing theory. III”, *Phys. Repts.* **203**, 125–231.
- de Boer, E. (1996). *The Cochlea*, chapter Mechanics of the Cochlea: Modeling Efforts, 258–317, Springer Handbook of Auditory Research (Springer-Verlag New York, Inc.).
- de Boer, E. and MacKay, R. (1980). “Reflections on reflections”, *J. Acoust. Soc. Am.* **67**, 882–890.
- de La Rochefoucauld, O., Kachroo, P., and Olson, E. S. (2010). “Ossicular motion related to middle ear transmission delay in gerbil.”, *Hear. Res.* **270**, 158–172.
- de La Rochefoucauld, O. and Olson, E. S. (2007). “The role of organ of Corti mass in passive cochlear tuning”, *Biophys. J.* **93**, 3434–3450.
- de Mestral, G. (1952). “Velvet type fabric and method of producing same”, United States Patent, no. 2,717,437, prior application in Switzerland, 1951.
- Diependaal, R. J., Duifhuis, H., Hoogstraten, H. W., and Viergever, M. A. (1987). “Numerical methods for solving one-dimensional cochlear models in the time domain”, *J. Acoust. Soc. Am.* **82**, 1655–1666.

- Diependaal, R. J. and Viergever, M. A. (1989). “Nonlinear and active two-dimensional cochlear models: time-domain solution.”, *J. Acoust. Soc. Am.* **85**, 803–812.
- Dong, W. and Olson, E. S. (2006). “Middle ear forward and reverse transmission in gerbil.”, *J. Neurophysiology* **95**, 2951–2961.
- Douglas, J. F., Gasiorek, J. M., and Swaffield, J. A. (2001). *Fluid Mechanics*, fourth edition (Pearson, Prentice Hall).
- Duifhuis, H. (2000). *Recent Developments in Auditory Mechanics*, chapter A general structure for a time domain model of the cochlea, 188–194 (World Scientific, Singapore).
- Duifhuis, H. (2011). “Hopf-bifurcations and Van der Pol oscillator model of the mammalian cochlea”, in *What Fire is in Mine Ears: Progress in Auditory Biomechanics*, 195–201 (American Institute of Physics, Melville, NY).
- Duifhuis, H. (2012). *Cochlear Mechanics Introduction to a Time Domain Analysis of the Nonlinear Cochlea* (Springer US).
- Duifhuis, H., Hoogstraten, H., van Netten, S., Diependaal, R., and Bialek, W. (1986). *Peripheral Auditory Mechanisms*, chapter Modelling the cochlear partition with coupled Van der Pol Oscillators, 290–297 (Springer, Berlin).
- Eguíluz, V. M., Ospeck, M., Choe, Y., Hudspeth, A. J., and Magnasco, M. O. (2000). “Essential nonlinearities in hearing.”, *Phys. Rev. Lett.* **84**, 5232–5235.
- Elliott, S. J., Ku, E. M., and Lineton, B. (2007). “A state space model for cochlear mechanics”, *J. Acoust. Soc. Am.* **122**, 2759–2771.
- Epp, B., Verhey, J. L., and Mauermann, M. (2010). “Modeling cochlear dynamics: Interrelation between cochlea mechanics and psychoacoustics.”, *J. Acoust. Soc. Am.* **128**, 1870–1883.
- Fawcett, D. W. (1986). *A Textbook of Histology*, twelfth edition (Chapman and Hall, New York).
- Fernández, C. (1952). “Dimensions of the cochlea (guinea pig)”, *J. Acoust. Soc. Am.* **24**, 519–523.
- Fish, J. and Belytschko, T. (2007). *A first course in finite elements* (Wiley).
- Flanagan, J. L. (1962). “Computational model for basilar-membrane displacement”, *J. Acoust. Soc. Am.* **34**, 1370–1376.
- Flax, M. R. (2008). “The active compression wave cochlear amplifier”, Ph.D. thesis, University of New South Wales.
- Fraginière, E. (1998). “Analogue VLSI emulation of the cochlea”, Ph.D. thesis, Lausanne, EPFL.

- Galambos, R. (2006). “Models and musings about them”, *Int. J. Psychophysiology* **60**, 101–105.
- Geisler, C. D. (1993). “A realizable cochlear model using feedback from motile outer hair cells”, *Hear. Res.* **68**, 253–262.
- Geisler, C. D. and Sang, C. (1995). “A cochlear model using feed-forward outer-hair-cell forces”, *Hear. Res.* **86**, 132–146.
- Givelberg, E. and Bunn, J. (2003). “A comprehensive three-dimensional model of the cochlea”, *J. Comp. Phys.* **191**, 377–391.
- Gold, T. (1948). “Hearing. II. the physical basis of the action of the cochlea”, *Proc. R. Soc. Lond.* **B135**, 492–498.
- Graham, P. (2004). *Hackers and Painters: Big Ideas from the Computer Age* (O’Reilly Media, Inc., 1005 Gravenstein Highway North, Sebastopol, CA, 95472.).
- Greenwood, D. D. (1961). “Critical bandwidth and the frequency coordinates of the basilar membrane”, *J. Acoust. Soc. Am.* **33**, 1344–1356.
- Greenwood, D. D. (1990). “A cochlear frequency-position function for several species—29 years later.”, *J. Acoust. Soc. Am.* **87**, 2592–2605.
- Gropp, W., Lusk, E., and Skjellum, A. (1999). *Using MPI, Volume 2*, second edition (MIT Press, Massachusetts).
- Gummer, A., Dalhoff, E., Nowotny, M., and Scherer, M., eds. (2003). *Biophysics of the Cochlea from Molecules to Models* (World Scientific Publishing Co. Pte. Ltd., 5 Toh Tuck Link, Singapore 596224).
- Hájek, O. (1979). “Discontinuous differential equations, I”, *J. Differential Eq.* **32**, 149–170.
- Hall, J. L. (1974). “Two-tone distortion products in a nonlinear model of the basilar membrane.”, *J. Acoust. Soc. Am.* **56**, 1818–1828.
- Hallworth, R. and Jensen-Smith, H. (2008). *Active Processes and Otoacoustic Emissions*, chapter The Morphological Specializations and Electromotility of the Mammalian Outer Hair Cell, 145–189 (Springer Science+Business Media, LLC).
- Hamilton, T. J. (2008). “Analogue VLSI implementations of two dimensional, nonlinear, active cochlea models”, Ph.D. thesis, University of Sydney, Sydney, Australia.
- Hamilton, T. J., Jin, C., van Schaik, A., and Tapson, J. (2008). “An active 2-d silicon cochlea”, *IEEE Transactions on Biomedical Circuits and Systems* **2**, 30–43.

- Hamilton, T. J., Tapson, J., Rapson, M., Jin, C., and van Schaik, A. (2010). “Understanding the mathematics of hearing using electronic circuits”, *ANZIAM J.* **51**, C300–C315.
- Hanekom, T. (2001). “Three-dimensional spiraling finite element model of the electronically stimulated cochlea”, *Ear & Hearing* **22**, 300–315.
- Helmholtz, H. (1954). *On the Sensations of Tone* (Dover Publications, Inc.), republication of the second (1885) edition of the Ellis translation of ‘Die Lehre von den Tonempfindungen’. Based on the fourth German edition of 1877.
- Henderson, H. V. and Searle, S. R. (1981). “On deriving the inverse of a sum of matrices”, *SIAM* **23**, 53–60.
- Heroux, M., Bartlett, R., Hoekstra, V. H. R., Hu, J., Kolda, T., Lehoucq, R., Long, K., Pawlowski, R., Phipps, E., Salinger, A., Thornquist, H., Tuminaro, R., Willenbring, J., and Williams, A. (2003). “An overview of Trilinos”, Technical Report SAND2003-2927, Sandia National Laboratories.
- Holley, M. C. (1996). *The Cochlea*, chapter Outer Hair Cell Motility, 386–434 (Springer-Verlag, New York).
- Hudspeth, A. J., Jülicher, F., and Martin, P. (2010). “A critique of the critical cochlea: Hopf–a bifurcation–is better than none”, *J. Neurophysiology* **104**, 1219–1229.
- Iserles, A. (2009). *A First Course in the Numerical Analysis of Differential Equations*, Cambridge Texts in Applied Mathematics, second edition (Cambridge University Press).
- Johannesma, P. I. M. (1980). “Narrow band filters and active resonators. Comments on papers by D. T. Kemp & R. A. Chum, and H. P. Wit & R. J. Ritsma”, in *Psychophysical, Physiological, and Behavioural Studies in Hearing*, edited by G. van den Brink and F. A. Bilten, 62–63 (Delft University Press, Delft).
- Johnstone, B. M. and Boyle, A. J. (1967). “Basilar membrane vibration examined with the mössbauer technique.”, *Science* **158**, 389–390.
- Kemp, D. T. (1978). “Stimulated acoustic emissions from within the human auditory system.”, *J. Acoust. Soc. Am.* **64**, 1386–1391.
- Kemp, D. T. (2008). *Active Processes and Otoacoustic Emissions*, chapter Otoacoustic Emissions: Concepts and Origins, 1–38, Springer Handbook of Auditory Research (Springer Science+Business Media, LLC).
- Kemp, D. T., Bray, P., Alexander, L., and Brown, A. M. (1986). “Acoustic emission cochleography—practical aspects.”, *Scand. Audiol. Suppl.* **25**, 71–95.

- Kemp, D. T. and Chum, R. (1980). “Properties of the generator of stimulated acoustic emissions.”, *Hear. Res.* **2**, 213–232.
- Kessel, R. G. and Kardon, R. H. (1979). *Tissues and Organs: a text-atlas of scanning electron microscopy* (W. H. Freeman, San Francisco).
- Kim, D. O., Molnar, C. E., and Matthews, J. W. (1980a). “Cochlear mechanics: nonlinear behavior in two-tone responses as reflected in cochlear-nerve-fiber responses and in ear-canal sound pressure.”, *J. Acoust. Soc. Am.* **67**, 1704–1721.
- Kim, D. O., Neely, S. T., Molnar, C. E., and Matthews, J. W. (1980b). *Psychological, Physiological and Behavioral Studies in Hearing*, chapter An active cochlear model with negative damping in the cochlear partition: comparison with Rhode’s ante- and post-mortem observations, 7–15 (Delft University Press, Delft, The Netherlands).
- Kim, Y. and Xin, J. (2005). “A two-dimensional nonlinear nonlocal feed-forward cochlear model and time domain computation of multitone interactions”, *Multiscale Model. Simul.* **4**, 664–690.
- Kleijn, W. B. (2008). *Springer Handbook of Speech Processing*, chapter Principles of Speech Coding, 283–305 (Springer-Verlag, Berlin Heidelberg).
- Kollmeier, B., Brand, T., and Meyer, B. (2008). *Springer Handbook of Speech Processing*, chapter Perception of Speech and Sound, 61–82 (Springer-Verlag, Berlin Heidelberg).
- Kolston, P. J. and Ashmore, J. F. (1996). “Finite element micromechanical modeling of the cochlea in three dimensions”, *J. Acoust. Soc. Am.* **99**, 455–467.
- Kros, C. J. (1996). *The Cochlea*, chapter Physiology of Mammalian Cochlear Hair Cells, 318–385 (Springer-Verlag, New York).
- Kros, C. J., Rüsçh, A., and Richardson, G. P. (1992). “Mechano-electrical transducer currents in hair cells of the cultured neonatal mouse cochlea.”, *Proc. Biol. Sci.* **249**, 185–193.
- Ku, E. M., Elliott, S. J., and Lineton, B. (2008). “Statistics of instabilities in a state space model of the human cochlea”, *J. Acoust. Soc. Am.* **124**, 1068–1079.
- Lange, K. (2004). *Optimization*, Springer Texts in Statistics (Springer-Verlag, New York).
- Lewis, R. S. and Hudspeth, A. J. (1983). “Voltage- and ion-dependent conductances in solitary vertebrate hair cells”, *Nature* **304**, 538–541.
- Lighthill, J. (1978). *Waves in Fluids* (Cambridge University Press, The Pitt Building, Trumpington Street, Cambridge CB2 1RP).

- Lighthill, J. (1981). “Energy flow in the cochlea”, *J. Fluid Mech.* **106**, 149–213.
- Lim, K.-M. and Steele, C. R. (2002). “A three-dimensional nonlinear active cochlear model analyzed by the wkb-numeric method.”, *Hear. Res.* **170**, 190–205.
- Liu, Y.-W. and Neely, S. T. (2010). “Distortion product emissions from a cochlear model with non-linear mechano-electrical transduction in outer hair cells.”, *J. Acoust. Soc. Am.* **127**, 2420–2432.
- Lüling, H., Franosch, J.-M. P., and van Hemmen, J. L. (2010). “A two-dimensional cochlear fluid model based on conformal mapping.”, *J. Acoust. Soc. Am.* **128**, 3577–3584.
- Lynch III, T. J., Nedzelitsky, V., and Peake, W. T. (1982). “Input impedance of the cochlea in cat”, *J. Acoust. Soc. Am.* **72**, 108–130.
- Lyon, R. F., Katsiamis, A. G., and Drakakis, E. M. (2011). “History and future of auditory filter models”, in *Proceedings of 2010 IEEE International Symposium on Circuits and Systems (ISCAS)*, 3809–3812.
- Lyon, R. F. and Mead, C. (1988a). “An analog electronic cochlea”, *IEEE Trans-ASSP* **36**, 1119–1134.
- Lyon, R. F. and Mead, C. A. (1988b). “Cochlear Hydrodynamics Demystified”, Technical Report Caltech-CS-TR-88-4, California Institute of Technology.
- MacFarlane, A. G. J. (1970). *Dynamical System Models* (George G. Harrap & Co. Ltd, 182 High Holborn, London, W.C.1).
- Mammano, F. and Nobili, R. (1993). “Biophysics of the cochlea: Linear approximation”, *J. Acoust. Soc. Am.* **93**, 3320–3332.
- Martin, P. (2008). *Active Processes and Otoacoustic Emissions*, chapter Active Hair-Bundle Motility of the Hair Cells of Vestibular and Auditory Organs, 93–143 (Springer Science+Business Media, LLC).
- Matthews, J. W. (1980). “Mechanical modeling of nonlinear phenomena observed in the peripheral auditory system”, Ph.D. thesis, Washington University, St. Louis, MO.
- Matthews, J. W. and Molnar, C. E. (1989). “On the stability of cochlear mechanical models”, in *Cochlear Mechanisms - Structure, Function and Models*, edited by J. P. Wilson and D. T. Kemp, 437–444 (Plenum Press).
- Meaud, J. and Grosh, K. (2010). “The effect of tectorial membrane and basilar membrane longitudinal coupling in cochlear mechanics.”, *J. Acoust. Soc. Am.* **127**, 1411–1421.

- Merchant, S. N. and Rosowski, J. J. (2008). “Conductive hearing loss caused by third-window lesions of the inner ear”, *Otol. Neurotol.* **29**, 282–289.
- Moleti, A., Paternoster, N., Bertaccini, D., Sisto, R., and Sanjust, F. (2009). “Otoacoustic emissions in time-domain solutions of nonlinear non-local cochlear models.”, *J. Acoust. Soc. Am.* **126**, 2425–2436.
- Mountain, D. C., Hubbard, A. E., Ketten, D. R., and O’Malley, J. T. (2003). “The helicotrema: Measurements and models”, in *Biophysics of the Cochlea: From Molecules to Models*, edited by M. N. A. W. Gummer, E. Dalhoff and M. P. Scherer (World Scientific, Singapore).
- Neely, S. T. (1978). “Mathematical models of the mechanics of the cochlea”, Master’s thesis, California Institute of Technology, Pasadena, California.
- Neely, S. T. (1981a). “Finite difference solution of a two-dimensional mathematical model of the cochlea”, *J. Acoust. Soc. Am.* **69**, 1386–1393.
- Neely, S. T. (1981b). “Fourth-order partition dynamics for a two-dimensional model of the cochlea”, Ph.D. thesis, Washington University, Saint Louis, Missouri.
- Neely, S. T. and Allen, J. B. (2009). *Proceedings of the 10th Mechanics of Hearing Workshop*, chapter Retrograde waves in the cochlea, 62–67 (World Scientific Publishing Co.).
- Neely, S. T. and Kim, D. O. (1986). “A model for active elements in cochlear biomechanics”, *J. Acoust. Soc. Am.* **79**, 1472–1480.
- O’Connor, K. N. and Puria, S. (2008). “Middle-ear circuit model parameters based on a population of human ears.”, *J. Acoust. Soc. Am.* **123**, 197–211.
- Olson, E. S. (1998). “Observing middle and inner ear mechanics with novel intracochlear pressure sensors.”, *J. Acoust. Soc. Am.* **103**, 3445–3463.
- Olukotun, K. and Hammond, L. (2005). “The future of microprocessors”, *Queue - Multiprocessors* **3**, 26–34.
- Parent, P. and Allen, J. B. (2007). “Wave model of the cat tympanic membrane.”, *J. Acoust. Soc. Am.* **122**, 918–931.
- Passchier-Vermeer, W. and Passchier, W. F. (2000). “Noise exposure and public health.”, *Environ. Health. Perspect.* **108 Suppl 1**, 123–131.
- Peterson, L. C. and Bogert, B. P. (1950). “A dynamical theory of the cochlea”, *J. Acoust. Soc. Am.* **22**, 369–381.

- Pickles, J. O. (1988). *An Introduction to the Physiology of Hearing*, second edition (Academic Press).
- Prell, C. G. L., Yamashita, D., Minami, S. B., Yamasoba, T., and Miller, J. M. (2007). “Mechanisms of noise-induced hearing loss indicate multiple methods of prevention.”, *Hear. Res.* **226**, 22–43.
- Price, K. V., Storm, R. M., and Lampinen, J. A. (2005). *Differential Evolution*, Natural Computing Series (Springer-Verlag, Berlin Heidelberg).
- Purchase, I. F. H. (2002). “Ethical issues for bioscientists in the new millennium.”, *Toxicol. Lett.* **127**, 307–313.
- Puria, S. and Allen, J. B. (1991). “A parametric study of cochlear input impedance”, *J. Acoust. Soc. Am.* **89**, 287–309.
- Puria, S. and Allen, J. B. (1998). “Measurements and model of the cat middle ear: Evidence of tympanic membrane acoustic delay”, *J. Acoust. Soc. Am.* **104**, 3463–3481.
- Rabiner, L. and Juang, B.-H. (1993). *Fundamentals of Speech Recognition* (Prentice Hall PTR).
- Ramamoorthy, S., Deo, N. V., and Grosh, K. (2007). “A mechano-electro-acoustical model for the cochlea: response to acoustic stimuli”, *J. Acoust. Soc. Am.* **121**, 2758–2773.
- Ramamoorthy, S., Zha, D.-J., and Nuttall, A. L. (2010). “The biophysical origin of traveling-wave dispersion in the cochlea”, *Biophys. J.* **99**, 1687–1695.
- Rapson, M. J., Hamilton, T. J., and Tapson, J. C. (2011). “Exploring the traveling wave hypothesis with a finite element, state space, time domain cochlear model”, EMAC 2011.
- Rapson, M. J., Karpul, D., and Tapson, J. C. (2012). “Unification and extension of monolithic state space and iterative cochlear models”, *J. Acoust. Soc. Am.* **131**, 3935–3952.
- Rapson, M. J. and Tapson, J. C. (2010). “Capturing low-frequency cochlear impedance in time domain model: The role of viscosity”, *J. Acoust. Soc. Am.* **127**, 1869–1869.
- Rapson, M. J. and Tapson, J. C. (2011a). “Emulating the theoretical bifurcation diagram of a cochlear model in computation”, AfriCOMP11.
- Rapson, M. J. and Tapson, J. C. (2011b). “Investigations into time stepping methods for cochlear models”, Appendix D in this work (originally submitted to *Int. J. Numer. Meth. Biomed. Engng*).
- Ren, T. (2002). “Longitudinal pattern of basilar membrane vibration in the sensitive cochlea”, *Proc. Nat. Acad. Sci.* **99**, 17101–17106.

- Rhode, W. S. (1971). "Observations of the vibration of the basilar membrane in squirrel monkeys using the Mössbauer technique.", *J. Acoust. Soc. Am.* **49**, Suppl 2:1218+.
- Rhode, W. S. (1973). *Basic Mechanisms in Hearing*, chapter An investigation of post-mortem cochlear mechanics using the Mössbauer effect, 49–67 (Academic Press, New York).
- Rhode, W. S. (1977). *Psychophysics and Physiology of Hearing*, chapter Some observations on two-tone interaction measured with the Mössbauer effect., 27–41 (Academic Press, London).
- Rhode, W. S. and Robles, L. (1974). "Evidence from Mössbauer experiments for nonlinear vibration in the cochlea", *J. Acoust. Soc. Am.* **55**, 588–596.
- Robles, L. and Ruggero, M. A. (2001). "Mechanics of the mammalian cochlea", *Physiological Reviews* **81**, 1305–1352.
- Robles, L., Ruggero, M. A., and Rich, N. C. (1991). "Two-tone distortion in the basilar membrane of the cochlea.", *Nature* **349**, 413–414.
- Ruggero, M. A. and Rich, N. C. (1987). "Timing of spike initiation in cochlear afferents: Dependence on site of innervation", *J. Neurophysiology* **58**, 379–403.
- Ruggero, M. A., Rich, N. C., Recio, A., Narayan, S. S., and Robles, L. (1997). "Basilar-membrane responses to tones at the base of the chinchilla cochlea.", *J. Acoust. Soc. Am.* **101**, 2151–2163.
- Ruggero, M. A., Robles, L., and Rich, N. C. (1992). "Two-tone suppression in the basilar membrane of the cochlea: mechanical basis of auditory-nerve rate suppression.", *J. Neurophysiology* **68**, 1087–1099.
- Russell, I. J. and Sellick, P. M. (1983). "Low-frequency characteristics of intracellularly recorded receptor potentials in guinea-pig cochlear hair cells.", *J. Physiology* **338**, 179–206.
- Russell, W. M. S. and Burch, R. L. (1959). *The Principles of Humane Experimental Technique* (Methuen, London).
- Ryan, A. F. and Dallos, P. (1984). *Physiology of the Cochlea*, chapter 22, 253–266, second edition (Little Brown, Boston).
- Salt, A. N. and Ma, Y. (2001). "Quantification of solute entry into cochlear perilymph through the round window membrane.", *Hear. Res.* **154**, 88–97.
- Sato, M., Leake, P. A., and Hradek, G. T. (1999). "Postnatal development of the organ of corti in cats: a light microscopic morphometric study", *Hear. Res.* **127**, 1–13.
- Schein, J. D. (2001). *Ethics in Deaf Education*, chapter Ethical Considerations in the Demography of Deafness, 21–32 (Academic Press).

- Senge, P. M. (1990). *The fifth discipline: the art and practice of the learning organization* (Doubleday/Currency, New York).
- Shera, C. A., Guinan, J. J., and Oxenham, A. J. (2002). “Revised estimates of human cochlear tuning from otoacoustic and behavioral measurements”, *Proc. Nat. Acad. Sc. USA* **99**, 3318–3323.
- Shera, C. A., Guinan, J. J., and Oxenham, A. J. (2010). “Otoacoustic estimation of cochlear tuning: Validation in the chinchilla”, *JARO* **11**, 343–365.
- Sisto, R., Moleti, A., Paternoster, N., Botti, T., and Bertaccini, D. (2010). “Different models of the active cochlea, and how to implement them in the state-space formalism.”, *J. Acoust. Soc. Am.* **128**, 1191–1202.
- Slepecky, N. B. (1996). *The Cochlea*, chapter Structure of the Mammalian Cochlea, 44–129, Springer Handbook of Auditory Research (Springer-Verlag, New York).
- Squillacote, A. H. (2007). *The ParaView guide: A parallel visualization application* (Kitware).
- Steele, C. R. and Taber, L. A. (1979a). “Comparison of wkb and finite difference calculations for a two-dimensional cochlear model.”, *J. Acoust. Soc. Am.* **65**, 1001–1006.
- Steele, C. R. and Taber, L. A. (1979b). “Comparison of WKB calculations and experimental results for three-dimensional cochlear models”, *J. Acoust. Soc. Am.* **65**, 1007–1018.
- Stenfelt, S., Hato, N., and Goode, R. L. (2004a). “Fluid volume displacement at the oval and round windows with air and bone conduction stimulation”, *J. Acoust. Soc. Am.* **115**, 797–812.
- Stenfelt, S., Hato, N., and Goode, R. L. (2004b). “Round window membrane motion with air conduction and bone conduction stimulation”, *Hearing. Res.* **198**, 10–24.
- Strang, G. (2007). *Computational Science and Engineering* (Wellesley-Cambridge Press, Box 812060, Wellesly MA 02482 USA).
- Stremmer, F. G. (1990). *Introduction to Communication Systems*, Addison-Wesley Series in Electrical Engineering, third edition edition (Addison-Wesley Publishing Company, Inc.).
- Strogatz, S. H. (1994). *Nonlinear dynamics and chaos: with applications to physics, biology, chemistry, and engineering* (Westview Press, Cambridge, MA).
- Sutter, H. and Larus, J. (2005). “Software and the concurrency revolution”, *Queue - Multiprocessors* **3**, 54–62.
- Thorne, M., Salt, A. N., DeMott, J. E., Henson, M. M., Henson, O. W., and Gewalt, S. L. (1999). “Cochlear fluid space dimensions for six species derived from reconstructions of three-dimensional magnetic resonance images”, *The Laryngoscope* **109**, 1661–1668.

- Tonndorf, J. and Tabor, J. R. (1962). “Closure of the cochlear windows: Its effect upon air- and bone-conduction”, *Ann. Otol. Rhinol. Laryngol.* **71**, 5–29.
- Ulfendahl, M. (1997). “Mechanical responses of the mammalian cochlea.”, *Prog Neurobiol* **53**, 331–380.
- van Hemmen, J. L., Kistler, W. M., and Thomas, E. G. F. (2000). “Calculation of Volterra kernels for solutions of nonlinear differential equations”, *SIAM J. Appl. Math.* **61**, 1–21.
- Viergever, M. (1980). “Mechanics of the inner ear”, Ph.D. thesis, Delft University of Technology.
- Vollrath, F. and Knight, D. P. (2001). “Liquid crystalline spinning of spider silk”, *Nature* **410**, 541–548.
- von Békésy, G. (1960). *Experiments in Hearing* (McGraw-Hill, New York).
- Vorländer, M. (2008). *Auralization*, 1st edition (Springer-Verlag Berlin Heidelberg).
- Warr, W. B. (1992). *Mammalian Auditory Pathway: Neuroanatomy*, chapter Organization of olivocochlear efferent systems in mammals, 410–448 (Springer-Verlag, New York).
- Watts, L. (1993). “Cochlear mechanics: Analysis and analog VLSI”, Ph.D. thesis, California Institute of Technology, Pasadena, California.
- Wegel, R. L. and Lane, C. E. (1924). “The Auditory Masking of One Pure Tone by Another and its Probable Relation to the Dynamics of the Inner Ear”, *Phys. Rev.* **23**, 266–285.
- Wever, E. G. (1949). *Theory of Hearing* (John Wiley and Sons, Inc., New York).
- WHO (2010). “Deafness and hearing impairment”, World Health Organization Fact Sheet No 300, URL <http://www.who.int/mediacentre/factsheets/fs300/en/index.html>, last seen 02/12/2010.
- Wilkinson, G. and Gray, A. A. (1924). *The Mechanism of the Cochlea: A Restatement of the Resonance Theory of Hearing* (Macmillan).
- Wolmarans, H. P. (2005). “Cochlear implant speech processing, based on the cochlear travelling wave”, Master’s thesis, University of Pretoria, Pretoria.
- Xin, J., Qi, Y., and Deng, L. (2002). “Computation and analysis of a nonlinear nonlocal cochlear model with applications to multitone interactions in hearing”, *J. Acoust. Soc. Am.* **112**, S2300–S2300, conference.
- Zweig, G. (1990). “The impedance of the organ of Corti”, in *The Mechanics and Biophysics of Hearing*, edited by P. Dallos, C. D. Geisler, J. W. Matthews, M. A. Ruggero, and C. R. Steele, Lecture Notes in Biomathematics, 362–369 (Springer-Verlag).

Zwislocki, J. J. (1948). "Theorie de Schneckenmechanik: Qualitative und Quantitative Analyse",
Acta Otolaryngol. **Suppl. 72.**, 1–76.